

NBER WORKING PAPER SERIES

WEATHERING THE STORM:
SUPPLY CHAINS AND CLIMATE RISK

Juanma Castro-Vincenzi
Gaurav Khanna
Nicolas Morales
Nitya Pandalai-Nayar

Working Paper 32218
<http://www.nber.org/papers/w32218>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2024, Revised March 2025

We thank Claire Conzelmann and Simon Farbman for outstanding research assistance and Swapnika Rachapalli for sharing data. We also thank our discussants David Atkin, Joaquin Blaum, Meredith Startz, Jose Vasquez, as well as Pol Antràs, Adrien Bilal, Chris Boehm, Olivier Coibion, Elhanan Helpman, Gabriel Kreindler, Andrei Levchenko, Hugo Lhuillier, Ezra Oberfield, Andres Rodriguez-Clare, Esteban Rossi-Hansberg, Alireza Tahbaz-Salehi, audiences at Minnesota Macro, Princeton IES, Stanford, University of Michigan, St Louis Fed, Warwick, McGill, UQAM, UT Austin, Columbia, CREI, NBER Summer Institute ITM, Midwest Macro, SED-Barcelona, WIEN - Vienna, TIGN, Empirical Macro Workshop and Kiel-CEPR Geopolitics Conference for helpful comments. The views expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Bank of Richmond, the Board of Governors, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2024 by Juanma Castro-Vincenzi, Gaurav Khanna, Nicolas Morales, and Nitya Pandalai-Nayar. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Weathering the Storm: Supply Chains and Climate Risk

Juanma Castro-Vincenzi, Gaurav Khanna, Nicolas Morales, and Nitya Pandalai-Nayar

NBER Working Paper No. 32218

March 2024, Revised March 2025

JEL No. E0, F0, F10, F18, F40, F62

ABSTRACT

We characterize how firms structure supply chains under climate risk. Using new data on the universe of firm-to-firm transactions from an Indian state, we show that firms diversify sourcing locations, and that suppliers exposed to climate risk charge lower prices. We develop a general equilibrium spatial model of firm input sourcing under climate risk. Firms diversify identical inputs from suppliers across space, trading off the probability of climate disruptions against higher input costs. We quantify the model using data on 271 Indian regions. Wages are inversely correlated with sourcing risk, giving rise to a cost minimization-resilience tradeoff. Supply chain diversification unambiguously reduces real wage volatility, but ambiguously affects their levels, as diversification may come with high input costs. While diversification mitigates climate risk, it exacerbates the distributional consequences of climate change by reducing wages in regions prone to frequent shocks.

Juanma Castro-Vincenzi
The Kenneth C. Griffin Department of Economics
University of Chicago
5757 S. University Avenue
Chicago, Illi 60637
castrovincenzi@uchicago.edu

Gaurav Khanna
School of Global Policy and Strategy
University of California, San Diego
9500 Gilman Drive
La Jolla, Cali 92093
United States
gakhanna@ucsd.edu

Nicolas Morales
Federal Reserve Bank of Richmond
701 E Byrd St
Richmond, VA 23219
nicolas.morales.uy@gmail.com

Nitya Pandalai-Nayar
Department of Economics
University of Texas at Austin
3.142 BRB, 2225 Speedway
Austin, TX 78713
and NBER
npnayar@utexas.edu

1 Introduction

The intersection of complex supply chains and climate risk presents a critical challenge to the global economy. Complex supply chains yield significant efficiency gains, enabling firms to procure inputs from the most efficient suppliers regardless of location. Yet, escalating climate risk raises concerns about the vulnerability of inter-linked production networks and the resultant broader economic fragility (Barrot and Sauvagnat, 2016; Boehm et al., 2019). Increasing climate change risk globally heightens the likelihood of natural disasters such as flooding and storm surges. In response, forward-looking firms might mitigate the impact of production disruptions through production location choices (Castro-Vincenzi, 2024) or supplier location diversification based on geographic variability in climate threats. Therefore, our understanding of how climate change may reshape economic production and its implications for welfare across regions hinges on firms' adaptive sourcing decisions. In this paper, we provide a theoretical, empirical, and quantitative analysis of the spatial consequences of supply chain restructuring in light of increased climate risk.

Studying the general equilibrium consequences of how firms structure supply chains when faced with climate hazards raises two important challenges. First, for empirical evidence on how firms adapt to climate risk, we need high-frequency data on transactions along the supply chain, the precise locations of establishments, and meaningful variation in weather-related events. Second, to quantify the broader economy-wide consequences, we require a general equilibrium model of firm input sourcing under climate risk, where firms face trade-offs such as the lower probability of climate shocks against higher costs, less productive inputs, or higher shipping costs.

To address the first issue, we obtain the universe of establishment-to-establishment level transactions from a large state in India, as long as one node of the transaction (buyer or seller) is in the state (the other node can be anywhere in the country). The data contains the precise establishment zip code, the value of the transaction, product code, date, quantity (and so the unit values), and the unique tax ID of the establishment. Using these data, we document important new motivating facts suggesting firms are optimizing supply chains to mitigate climate risk. First, firms diversify the locations they source from, multisourcing 74% of product value even within narrow product codes. Second, firms that multisource the same product buy from farther distances and from drier regions and pay higher prices. And third, suppliers in regions

that are more exposed to climate risk tend to charge lower prices.

An advantage of our setting is that India experiences monsoonal rainfall that follows a somewhat predictable spatial pattern every year, although the intensity and timing can vary. Regions across India regularly experience large flooding events that disrupt firm supply chains. Firms operating in this environment might reasonably consider the probability of climate-related disruptions in their operations, as suggested by our descriptive analysis.

To provide causal evidence of firm responses to climate shocks, we leverage the exogenous geographic and temporal variation in flooding events using event study designs. We show that the sales of flood-hit suppliers fall sharply over three months but recover by five months after a flood. The total purchases and sales of downstream buyers decrease substantially. Firms recover relatively quickly, and are unlikely to substitute to other suppliers, in contrast to the supply-chain reorganization documented by [Khanna et al. \(2025\)](#) following the unanticipated COVID-19 lockdowns. Our descriptive and event-study results are suggestive that firms plan for climate risk and face an input-cost vs disruption risk trade-off in setting up supply chains.

Our second contribution is theoretical. To address the challenge of quantifying economy-wide impacts, we build a new spatial general equilibrium model of firm sourcing under risk. Motivated by our empirical evidence, firms diversify their sourcing of otherwise-identical inputs across locations to mitigate climate risk. Such diversification comes with a trade-off: in general equilibrium, input prices are higher for places with lower climate risk, which might also be less productive or geographically distant, necessitating payment of higher trade costs.

A key feature of the model is that firms' expected profit functions in the presence of sourcing risk are concave in input orders. That is, firms behave as if they are risk-averse, even in the absence of explicit managerial risk aversion.¹ This implies that firms from each region will choose to diversify their input sourcing across regions if they face imperfectly correlated regional risk, even if regional fundamentals are constant across space and trade is costly (a "symmetric" economy). In a comparative statics exercise, we show that in this setting, where there is no love-for-variety trade motive, trade still occurs due to the diversification motives of firms. As a result, despite identical fundamentals, "safer" regions see higher real wages in general

¹[Blaum et al. \(2024\)](#) study firm input sourcing under shipping time risk where firms face a similar problem. In contrast, our focus is the multi-region general equilibrium.

equilibrium, while riskier regions see lower real wages.

Interestingly, this comparative statics exercise implies that the prices of inputs, and therefore, of regional consumption, are higher under costly trade than autarky. A stark insight from this exercise is that expected regional real wages can be lower under costly trade than under autarky, but their volatility is also lower. With risk aversion in consumer preferences, the decrease in volatility offsets the decline in expected real wages, and diversification is welfare-improving, but aggregate output is lower.

We quantify this model using a census of manufacturing firms across the country, allowing us to estimate location-specific productivities and labor shares. We implement the model on 271 regions in India. To discipline the magnitude of disruptions, we leverage the estimated input disruptions from our event studies.

Our model implies that bilateral sourcing shares are a function of all regional labor endowments, productivities, and bilateral trade costs, as well as the risk of sourcing in each region. Given estimates of regional labor, productivity, and bilateral trade costs, we back out the model-implied spatially correlated regional risk to match observed sourcing shares. To validate our framework, we project the model-implied risk on climate observables such as rainfall, floods, temperature and dryness as well as other risk-related variables such as state fixed effects, ruggedness, and elevation to capture institutional and geographic features that affect firm decisions.² We find that climate-related risk is strongly positively correlated with the estimated risk probabilities, with an R^2 of 0.32. While not causal, the robust positive correlation is consistent with firms taking into account several sources of risk when they form their supply chains, a feature that has been largely ignored by the literature (an exception is [Kopytov et al. \(2024\)](#) who study how supply chains adapt to supplier volatility).³

In contrast to the comparative statics, the effects of firm diversification in a realistic economy will depend on the variation in fundamentals in addition to risk-mitigation incentives. Our third contribution is, therefore, quantitative: we compute expected real wages, real wage volatility, and welfare across districts in our calibrated model, given model-implied sourcing risk. Our framework implies that as a result of firm

²While these parameters in the model are estimated conditional on productivity, it is well-known that cross-sectional climate risk, globally, is negatively correlated with productivity. To mitigate confounding, we also control for regional productivity in the projections.

³In an alternative exercise, we parameterize the regional risk as a function of climate-related variables and other risk-related variables related to institutional quality and local development and estimate the relevant parameters. Our quantitative results remain very similar.

sourcing decisions, real wages in each district will depend on the geography, productivity, and climate risk of all districts.⁴

We perform several quantitative exercises in our calibrated model. First, we quantify the insight from our comparative statics exercises regarding wage volatility. We find that under the estimated trade costs and climate risk, the variance of real wages is 9.25% higher in autarky. Expected real wages are also higher in autarky. In our calibrated model, they are, on average, 3.1% higher, although for some districts expected real wages decline. With log utility, autarky is welfare reducing, with a 7.29% average welfare decline.

We then study how regional wages change in general equilibrium under alternative shock probabilities to capture scenarios of changing climate risk, and to highlight our new channel. We use the correlation of flood, heat, dryness, and precipitation risk with our estimated district-level risk probabilities to infer how these probabilities would change given IPCC projections of climate risk. We then compute expected real wages, input prices, and wage volatility under the scenario of climate evolving as projected, holding all other long-run changes such as productivity growth constant. We find that the average risk of districts increases by 1.1p.p., but there is wide heterogeneity. Expected real wages decline on average by 1.96%, their volatility increases slightly by an average of 0.15%. Welfare decreases on average by 2.01%. Around 37% of districts see expected real wage increases.

Our quantification highlights the distributional consequences of adaptation to climate risk. In the counterfactual, initially better-off districts largely see welfare increases, while initially worse-off districts see welfare declines. We decompose the changes into the direct effects of changing risk and equilibrium effects of adaptation. Regions where risk is increasing bear the direct effects of shocks, but also see downward pressure on wages due to firms' adaptation away from them. We show that for regions that are experiencing an increase in risk, the economy's adaptation is welfare decreasing. In sum, our model and quantification show that firm sourcing decisions help mitigate the effects of climate shocks and have quantitatively important general equilibrium implications for real wages in safer regions relative to riskier ones.

Our results highlight two economic implications of climate change. On the positive

⁴For expositional simplicity, we use the term "risk" throughout the paper, but note that shocks in our model have mean and variance effects. In our quantification, we decompose the effects of risk into first moment (expected real wage) and second moment (volatility) effects.

side, the risks of climate change are partially mitigated as firms anticipate climate risk and diversify their sourcing decisions. This implies that more volatile weather does not necessarily translate into higher aggregate output volatility. On the negative side, climate change will have even larger redistributive effects across regions than commonly believed. Regions with more climate risk will face the direct effects of the shocks themselves, but additionally will also become less appealing to other regions as a source of inputs. As a result, demand for products from these regions will decline, and real wages will fall. The converse will occur in “safer” regions. In other words, diversification amplifies the distributional effects of climate change.

Related literature. A growing literature studies how climate change shapes economic activity in the long run, assessing how the distribution of economic activity changes within and across regions, and countries (Desmet et al., 2021, Jia et al., 2022, Cruz and Rossi-Hansberg, 2024, Hsiao, 2023, Bilal and Rossi-Hansberg, 2023, Balboni, 2025, Farrokhi and Lashkaripour, 2024, Nath, 2024). Another branch of the literature studies the effects of extreme weather events on firms’ employment and location decisions, as well as on FDI, using empirical studies or stylized theories (Indaco et al., 2020, Gu and Hale, 2023, Pankratz and Schiller, 2024). Both this paper and Castro-Vincenzi (2024) examine how changes in disruption probabilities from extreme weather events shape firms’ investments to mitigate risks—this paper through supplier diversification and Castro-Vincenzi (2024) via plant relocation. However, Castro-Vincenzi (2024) focuses on modeling in detail the industry equilibrium of the global car industry, whereas this paper solves for the full general equilibrium of a spatial economy under any distribution of location-specific risk.

Our theoretical and quantitative results are related to Kopytov et al. (2024), who study supply chain adaptation to supplier volatility, and to Pellet and Tahbaz-Salehi (2023), who study the implications of rigidities in supply chains that arise due to incomplete information. Similar to the rigid inputs in Pellet and Tahbaz-Salehi (2023), firms in our model place orders for intermediate inputs prior to shock realization, and cannot adjust orders ex-post. In contrast to these papers, our model features households in multiple regions who cannot trade shares of the different firms, and the incentive to mitigate volatility arises from the concavity of firm profits. As a result, in our framework, aggregate volatility decreases in trade openness, as firms mitigate risk, reminiscent of the findings in Caselli et al. (2019). However, expected real wages can be lower under costly trade compared to autarky. This parallels the

results in these papers that aggregate output is also lower due to diversification away from volatile suppliers. In our setting, eliminating trade barriers permits both expected real wages to be higher and aggregate volatility to be lower, maximizing the benefits of diversification. At the micro level, our firm problem is similar to [Blaum et al. \(2024\)](#), but our model delivers strong implications for how wages across space are shaped by regional risk in general equilibrium, and can be used to infer the risk that firms assign to different sourcing locations.

Supply chain fragility and resilience have received increased attention in the literature following recent global events ([Grossman et al., 2023, 2024](#); [Khanna et al., 2025](#); [Korovkin et al., 2024](#)). Our firm-to-firm data are similar to [Khanna et al. \(2025\)](#), but our identification strategy uses extreme weather events, and we emphasize the general equilibrium consequences of the adaptation of supply chains to climate risk, which are not studied in that paper. Our empirical evidence indeed suggests that firms' supply chain responses to climate-related risk vary qualitatively and quantitatively from their responses to an unanticipated, temporary shock like COVID-19.

A large research agenda emphasizes the importance of international trade in inputs and the macroeconomic consequences of such trade ([Yi, 2003](#), [Johnson and Noguera, 2012](#), [Caliendo and Parro, 2015](#), [Antràs et al., 2017](#), [Huo et al., 2024](#)). Some papers study the transmission of natural disasters through trade and supply chain links ([Barrot and Sauvagnat, 2016](#); [Boehm et al., 2019](#); [Carvalho et al., 2021](#)). Our focus is on quantifying the general equilibrium economy-wide consequences of firm supply chain adaption to the (changing) probability of disruptions, rather than firm responses to the incidence of a disruption.

Finally, our paper contributes to research studying trade under risk (e.g. [Helpman and Razin, 1978](#), [Esposito, 2022](#), [Allen and Atkin, 2022](#), [Adamopoulos and Leibovici, 2024](#), among others). [Balboni et al. \(2024\)](#) and [Blaum et al. \(2024\)](#) provide evidence of firm adaptation in Pakistan and the US, respectively. These papers provide empirical evidence which we complement, but we focus on the quantitative model studying the spatial general equilibrium implications of supply chain adaptation to risk.

The rest of our paper is structured as follows. Section 2 outlines our data and shows descriptive patterns. Section 3 sets up the model, derives some analytical results, and performs comparative statics. Section 4 calibrates and quantifies the model and contains the climate change counterfactuals. Section 5 concludes.

2 Empirical Approach

2.1 Data

Firm-to-firm trade. Our primary data source is daily establishment-level transactions (while we use the term “firm”, the data are at the granular establishment level). These data are from the tax authority of a large Indian state with a fairly diversified production structure, roughly 50% urbanization, and high population density. Comparing this context to others with firm-to-firm transaction data, the state has roughly three times Belgium’s population, seven times Costa Rica’s, and double Chile’s.

The data contain daily transactions from April 2018 to October 2020 between all registered establishments within the state, and all transactions where one node of the transaction (either buyer or seller) is in the state. All transactions have unique tax identifiers for both the selling and buying establishments, and we observe the value of the whole transaction, the value of the items being traded by 8-digit HSN code, the quantity of each item, its unit, and transportation mode.

Each transaction also reports the zip code location of both the selling and buying establishments, which we merge with other geographic data. By law, any goods transaction with value over Rs.50,000 (\$700) has to generate away-bills, which populate our data. Transactions with values lower than \$700 can also optionally be registered. As such, our network is representative of relatively larger firms, but the threshold is sufficiently low to capture small firms too. Indeed, part of the switch away from a traditional VAT (value-added tax) to the Goods and Services Tax (GST) regime was to expand the tax base and include many smaller establishments. The tax base under this GST regime includes small (as small as one worker) and large establishments. More information is in Appendix A, with summary statistics in Table A1. The distribution of customers and suppliers of each firm is very similar to that documented by [Alfaro Ureña et al. \(2018\)](#) for Costa Rica.

We use the data to construct the buyer-supplier network every period, the total value of firms’ inputs purchased, and output sold. To obtain a measure of real inputs and output, we use the reported quantity of each transaction to calculate product unit values, construct price indices, and deflate firm-level input purchases and sales.

2.2 Descriptive Analysis

To begin, we document three facts related to supplier diversification and climate risk to motivate the key features of our model.

Fact 1: Many firms source the same product from multiple regions. We leverage the detailed product information in our transaction data and compute the number of districts from which a firm sources a given product. As shown in Table 1 Columns 1 and 2, 96.5% of the total value of purchases come from firms with more than one supplier-district, corresponding to 62% of firms purchasing from more than one district. In Columns 3 to 6, we show that a significant fraction of firms also multisource the same product across regions. We compute the number of districts a firm-by-HSN product code pair sources from. In Columns 3 and 4, we use 4-digit product codes; and in Column 5-6, 8-digit product codes. Even with the narrowest product definition available in our data, 14.4% of firms source the same product from more than one district, and 74% of purchases come from firms that source the product from more than one district. This is evidence that a significant fraction of firms multisource their products. In Appendix Table B1, we show that the distribution of the number of supplier-districts is very similar when we exclude likely wholesalers and likely retailers from the analysis.⁵

Table 1: Share of firms that source from multiple districts

Number of supplier districts	Share of buyers		Share of buyers x HSN-4		Share of buyers x HSN-8	
	Firms	Value	Firms	Value	Firms	Value
1	37.9%	3.5%	77.1%	13.1%	85.6%	25.8%
2-5	45.8%	20.1%	21.2%	39.2%	13.8%	42.1%
6-9	9.9%	14.4%	1.2%	17.2%	0.4%	13.5%
10+	6.4%	62.0%	0.4%	30.5%	0.1%	18.6%

Note. Column 1 aggregates the data at the firm level and computes the share of firms that source from a certain number of districts. Column 2 calculates the fraction of total value purchased by number of supplier districts sourced from. Columns 3-4 aggregate the data at the firm-by-4-digit product level, and Columns 5-6 at the firm-by-8-digit product level.

Fact 2: Firms that multisource more pay higher input prices, and also buy products from farther distances and dryer regions. Focusing on firm-product pairs at the 8-digit product level in Figure 1a, we show that firms that source the

⁵Table B2 shows that the results are consistent when looking at multisourcing across firms instead of supplier districts.

same product from more regions tend to buy from suppliers that are farther away. For instance, firm-product pairs that source from one district have an average distance of 350km to suppliers. On the other hand, firm-product pairs with five suppliers per product more than double average distance, at 711km.

Figure 1b shows that firm-product pairs with more suppliers also seem to source from less rainy districts. For firm-product pairs that source from one district, such districts have, on average, 6.5mm daily rainfall. On the other hand, for firm-product pairs that source from five districts, such districts have, on average, 5.4mm of daily rainfall. The 1.1mm difference in rainfall between one and five source districts is 17% with respect to the mean. In Appendix Table B3, we show that such patterns are also prevalent for other measures of climate risk, such as historical riverine flooding.

Finally, in Figure 1c, we show that firms that source from more districts also tend to pay higher prices for their inputs.⁶ As shown in Figure 1c, firms that source from five districts pay an average price that is almost one standard deviation higher than firms that source from only one district. The average price paid monotonically increases with the number of districts sourced from.⁷

Fact 3: Supplier districts that face higher climate risk charge lower prices.

Figures 1b and 1c suggest that as buyers purchase from more suppliers, they source from regions with lower climate risk and pay higher prices. The flip side of this pattern is that suppliers in riskier areas might charge lower prices. To investigate this relationship further, we estimate a regression at the buyer (j) - supplier district (d) - product (p) level as in equation 1.

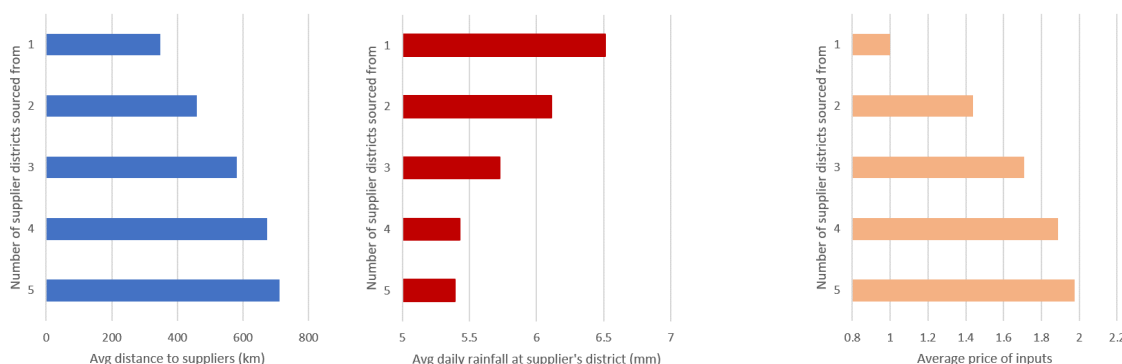
$$\log(\text{Price})_{j,d,p} = \alpha_1 \log(\text{Climate risk})_d + \alpha_2 \log(\text{Distance})_{j,d} + \alpha_3 \mathbb{1}(j \text{ in } d)_{j,d} + \alpha_4 \mathbb{1}(j, d \text{ in same state})_{j,d} + \gamma X_{d,p} + \delta_j + \delta_p + \epsilon_{j,d,p}, \quad (1)$$

where $\log(\text{Price})_{j,d,p}$ is the log of the average price charged to buyer j for product p by suppliers in district d . We control for the distance between j and d , indicators on whether the buyer is in district d or the same state as district d , and a set of controls at the product-supplier district level ($X_{d,p}$) such as the log size of all suppliers' sales from

⁶To compute average prices, we first estimate a regression of log price on product fixed effects, and standardize the residual of such regression to construct our residual price index. We then normalize the average price for those firms that source from only one region to one.

⁷In Appendix Table B3, we show that these patterns are statistically significant, and remain so within product and controlling for buyer size and supplier size. In other words, the patterns are not driven by specific products, by larger buyers, or by supplier capacity.

Figure 1: Supplier characteristics by number of districts sourced from



(a) Distance to suppliers (b) Average rainfall (c) Average price of inputs

Note. In the left panel, we compute the average distance between the firm and each of its suppliers from our transaction data. We then compute the average distance across firm-product pairs sourcing from 1 to 5 districts. In the middle panel, for each firm-product pair, we compute the average daily rainfall at each of the districts the firm sources from. Daily rainfall comes from the India Meteorological Department. We then compute the average across all firm-product pairs sourcing from 1 to 5 districts. In the right panel, we compute the average price paid for inputs for firm-product pairs sourcing from 1 to 5 districts. To construct our price index, we first run a regression of log prices on product fixed effects and take the residual. We standardize the residual and normalize it to 1 for firm-product pairs that source from only one district.

that district-product pair and the log of the total sales from that district. We also include buyer and product fixed effects, so the identification of the climate variables comes from firms that buy from multiple districts. Additionally, we include covariates that aim to capture market power at the supplier district, such as the log of the total number of suppliers for a given product in the district and the log of the largest supplier market share for that product in the district.

We consider two climate risk measures: the average daily rainfall for each district in 2019 and the historical river flooding in each district. Appendix B.2 details how these climate variables are computed. As shown in Table 2, both climate measures are negatively correlated with prices. The magnitudes are robust to including additional controls at the supplier-district level. A 10% increase in rainfall in a district is associated with suppliers in those districts charging 0.11% lower prices. Similarly, a 10% increase in riverine flooding levels in a district is associated with 2.55% lower prices charged by suppliers in that district. While these results cannot be interpreted as causal, they are suggestive that riskier areas charge lower prices.

Table 2: Correlation between price and supplier district climate risk

	Log (Price) _{<i>j,d,p</i>}	Log (Price) _{<i>j,d,p</i>}		Log (Price) _{<i>j,d,p</i>}	Log (Price) _{<i>j,d,p</i>}
Log(Avg Rainfall) _{<i>d</i>}	-0.0179*** (0.005)	-0.0112** (0.005)	Log(Avg River Flooding) _{<i>d</i>}	-0.381*** (0.026)	-0.255*** (0.026)
N obs	991,802	991,802	N obs	996,720	996,720
Additional controls	No	Yes	Additional controls	No	Yes

Note. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. We run a cross-sectional regression at the firm (j), supplier district (d), 8-digit product (p) level. The outcome is the log of the average price charged by suppliers in district d , to firm j for product p . The first and third columns control for log average distance between j and suppliers in d , a dummy variable for whether j is in district d , a dummy variable for whether j is in the same state as d , the log of total sales in product p from suppliers in d , the log of total sales of suppliers in d across all products, buyer fixed effects and product fixed effects. Columns 2 and 4 include controls for the log number of suppliers for product p in d and the log market share of the highest supplier of product p . Climate variables used are average daily rainfall in district in 2019 (left panel) and historical riverine flooding levels in district (right panel).

Fact 4: Purchases fall temporarily when suppliers are affected by floods.

Next, we leverage the timing of unexpected floods to examine how input purchases change in the lead-up to and right after the shock. Our event study examines pre-trends in the lead-up to the shock, and dynamics thereafter. The absence of pre-trends suggests that our parallel-trends identification assumption is likely to hold, whereas the post-shock dynamics are informative of how long it takes for firms to recover after the flood. In Appendix B.1, we discuss the data on flood events in our sample.

We use the existing supplier network (in the pre-shock period) as a measure of the exposure to the disruption to study how buyers were affected when their suppliers were hit. We examine outcomes $y_{j,t,k,\tau}$ for firm j , in period t , and industry k , measured in event-time (since flood) τ using the specification:

$$y_{j,t,k,\tau} = \sum_{x=-5}^{x=+5} [\gamma_x (\text{Flood Exposure})_{j\tau} + \delta_{\tau,x} + \beta_x X_{j,\tau_0-1}] + \delta_j + \delta_{r,k,t} + \epsilon_{j,t,k,\tau} \quad (2)$$

We estimate two specifications. First, when studying the direct impacts on suppliers in flooded areas, “Flood Exposure_{*j* τ} ” takes a value of 1 if firm j was exposed to a flood. Then, when examining how downstream buyers are affected, “Flood Exposure_{*j* τ} ” is “Supplier Exposure_{*j* τ} ”, capturing how exposed its suppliers were to the flood:

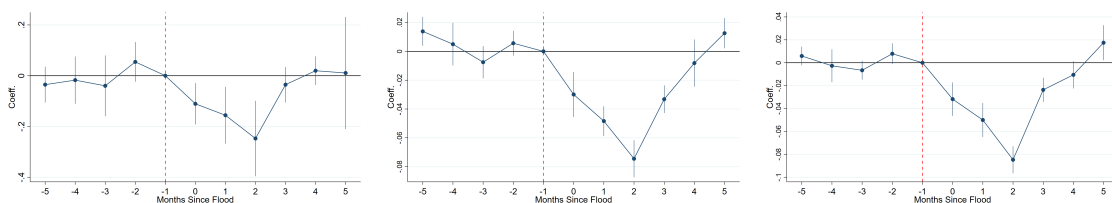
$$(\text{Supplier Exposure})_{j\tau} = \sum_i^N s_{i,j,\tau,x < 0} \times \mathbb{1}(\text{Supplier } i \text{ exposed to flood in } \tau) ,$$

where $s_{i,j,\tau,x < 0}$ is the value of purchases that firm j buys from firm i , relative to firm j 's total purchases, over the five months before the flood. The index essentially calculates the weighted average of the flood exposure of firm j 's sellers. A higher value of the index implies firm j faces a higher “supplier-exposure,” as a larger share of its purchases were coming from firms exposed to the flood.

We include a wide range of high-dimensional fixed effects to account for confounding shocks. Firm fixed effects δ_j control for firm-specific time-invariant differences; industry-by-time fixed effects $\delta_{r,k,t}$ control for district-industry-specific shocks and any demand shocks;⁸ and flood event-time since flood fixed effects $\delta_{\tau,x}$ control for aggregate trends around the flood event that affect all firms (including those not in the flood-exposed areas). X_{j,τ_0-1} contains controls for firm-demand shocks by including the pre-period exposure to floods of consumers, interacted with time indicators. It also includes controls for firm size-specific shocks by controlling for purchases in the pre-period, interacted with time-since flood indicators.

Figure 2a plots impacts on suppliers, and shows a lack of meaningful pre-trends in the lead-up to the flood. After the flood, there is an immediate decline in sales of 0.10 log points, which worsens until two months after the flood. After the two-month slump, there is a quick recovery to what they were in the pre-period.

Figure 2: Effects of Floods on Sales and Purchases



(a) Sales of affected suppliers (b) Downstream purchases (c) From returning suppliers

Note. Figure 2a includes event-time, industry-time, and firm fixed effects and controls for pre-period firm sales interacted with time indicators. Figure 2b and 2c include firm, time, event-time, and industry-district-real time fixed effects, and log pre-period purchases-time controls. We also include firm-demand controls by including the pre-period exposure to floods of a firm’s consumers interacted with time dummies. Standard errors clustered at the district level.

Figure 2b plots effects on purchases of downstream firms. Once again, the coefficients in the pre-periods do not display any meaningful trends. Consistent with Figure 2a, we find that purchases decline sharply for the first few months, and then

⁸When studying the direct impacts on suppliers, we include industry-time fixed effects $\delta_{k,t}$ instead, as the flood varies at the district-time level.

start to recover. Purchases are the lowest at two months after the flood, dropping by 0.07 log points with respect to the baseline period, for every one standard deviation increase in the supplier exposure (SD of exposure is 0.1). As we describe in section 4.2, we use this estimate to choose the input disruption parameter χ_j to match the drop. Interestingly, Figure 2c shows that affected firms return to existing suppliers (rather than switch suppliers), which may suggest that firms are adapting to the known risk of climate-related disruptions ex-ante.

Appendix B.1.2 examines a wider range of outcomes and methods. It shows these patterns are similar if we use insights from recent advances in two-way fixed effects methods, and estimate Local Projections-Diff-in-Diff (LP-DID) specifications or use a binary treatment. We also illustrate responses of other outcomes, such as downstream sales and prices.⁹ Finally, in Appendix C.4, we examine the importance of inventories in our empirical analysis. We find that inventories are, on average, less than a month’s sales, and are not correlated with multisourcing behavior.

In sum, our descriptive analysis, while not causal, provides suggestive evidence consistent with firms diversifying inputs to mitigate climate risk, and in the process facing a trade-off between input costs and risk.

3 Model

We develop a spatial general equilibrium model of firm sourcing under risk and perform comparative statics. The model is static, as rich geographic variation and a large number of locations is necessary for illustrating the diversification mechanism.¹⁰

⁹We find that downstream sales decrease by 7% in the three months after the shock for exposed firms relative to non-exposed firms while total purchases decrease by 16%, implying for every 1% decrease in purchases, sales decrease by 0.43%. In a back-of-the-envelope calculation, our quantitative model in Section 4 predicts that for every 1% decrease in purchases, sales decrease by 0.47%.

¹⁰The event studies illustrated (short-lived) dynamic responses to shock incidence. However, a dynamic model with sufficiently rich geographic variation is not currently tractable. Our emphasis is on understanding the steady state GE consequences of a distribution of risk across space, not on the reaction to the incidence of a disruption. That said, our model can be used to study the immediate ex-post response to the incidence of a disruption, as we do in Section 4.5. Given the short-lived responses in the data, we would not expect a significant role for additional dynamics here.

3.1 Setting

The economy consists of I regions. Each region i is endowed with L_i workers, a unit continuum of final goods producers who produce nontraded final goods, and competitive intermediate goods producers.

Timing. The model is static and consists of two stages. In the first stage, final goods producers in each location i place their orders for intermediate inputs from location j , M_{ji} . In the second stage, inputs are produced, origin-specific sourcing disruption shocks, $\boldsymbol{\chi} = \{\chi_j\}, j \in I$ are realized, and then inputs are delivered, final goods firms choose their labor inputs and produce, households supply labor and consume, and all markets clear at equilibrium prices.

Households. The representative household in region i supplies labor L_i inelastically to firms in i and chooses a consumption aggregate of the non-traded regional final goods, $q_i(\boldsymbol{\chi})$, to maximize

$$\max_{q_i(\omega, \boldsymbol{\chi})} \log \left(\left[\int_{\omega \in [0,1]} q_i(\omega, \boldsymbol{\chi})^{\frac{\sigma-1}{\sigma}} d\omega \right]^{\frac{\sigma}{\sigma-1}} \right) \quad (3)$$

subject to the budget constraint,

$$\int_{\omega \in [0,1]} p_i(\omega, \boldsymbol{\chi}) q_i(\omega, \boldsymbol{\chi}) = Y_i(\boldsymbol{\chi}) \equiv w_i(\boldsymbol{\chi}) L_i + \Pi_i(\boldsymbol{\chi}) \quad \forall \boldsymbol{\chi} \in \mathcal{G}(\boldsymbol{\chi}), \quad (4)$$

where $p_i(\omega, \boldsymbol{\chi})$ is the price of final good $q_i(\omega, \boldsymbol{\chi})$, $Y_i(\boldsymbol{\chi})$ is total income in region i , and $\sigma > 1$ is the elasticity of substitution. Total income $Y_i(\boldsymbol{\chi})$ is composed of labor income, $w_i(\boldsymbol{\chi}) L_i$, and aggregate profits rebated to the household by the firms, $\Pi_i(\boldsymbol{\chi})$. Our baseline model assumes labor is immobile across regions.

The Lagrange multipliers $\lambda_i(\boldsymbol{\chi}) = \frac{1}{Y_i(\boldsymbol{\chi})}$ of the state-specific budget constraints measure how much an extra unit of income contributes to utility in different states of the world. These multipliers define the stochastic discount factor firms use to compare profits across different states of the world.

Intermediate goods producers. In each region, there are a continuum of competitive suppliers of tradable intermediate inputs, \bar{M}_i , with production function $\bar{M}_i = z_i \ell_i^M$, where z_i is their productivity and ℓ_i^M is the labor used in the production of intermediates. The price of intermediates in i is equal to their constant marginal cost, $p_i^M(\boldsymbol{\chi}) = \frac{w_i(\boldsymbol{\chi})}{z_i}$, where $w_i(\boldsymbol{\chi})$ corresponds to the wages in that region. Notice

that intermediates \bar{M}_i are produced before the realization of shocks, but their price is potentially stochastic.

Let $p_{ji}^M(\boldsymbol{\chi})$ denote the price of intermediates from j used in i . We assume iceberg trade costs τ_{ji} between regions. No arbitrage in shipping implies that the factory-gate price and price at time of intermediate usage are related: $p_{ji}^M(\boldsymbol{\chi}) = \tau_{ji} p_j^M(\boldsymbol{\chi})$.

Final goods firms. Each region i contains a unit continuum of homogeneous final goods producers that produce differentiated varieties ω . Final goods are not tradable across regions. The constant returns to scale production function of the firms is

$$q_i(\omega, \boldsymbol{\chi}) = \phi_i \ell_i(\omega, \boldsymbol{\chi})^\beta x_i(\omega, \boldsymbol{\chi})^{1-\beta}, \quad (5)$$

where ϕ_i is the productivity of final goods' producers in location i , $\ell_i(\omega, \boldsymbol{\chi})$ is the firm's labor input, and intermediates, $x_i(\omega, \boldsymbol{\chi})$, can be sourced from each region $j \in I$ as perfect substitutes:¹¹

$$x_i(\omega, \boldsymbol{\chi}) = \sum_{j \in I} x_{ji}(\omega, \boldsymbol{\chi}). \quad (6)$$

For compact notation, for the remainder of the paper we suppress the explicit dependence of variables on $\boldsymbol{\chi}$ except where necessary for expositional clarity. Note that all equilibrium variables *except* \bar{M}_i remain potentially stochastic.

Second stage. In the second stage, final goods firms have already placed their orders of intermediates $M_{ji}(\omega)$, shocks $\boldsymbol{\chi}$ have been realized, and production takes place. The second-stage profit maximization problem of a final goods firm in i is

$$\max_{q_i, \{x_{ji}\}_{j=1}^I, \ell_i} [Y_i \mathbb{P}_i^{\sigma-1}]^{\frac{1}{\sigma}} q_i(\omega)^{\frac{\sigma-1}{\sigma}} - w_i \ell_i(\omega) \quad (7)$$

$$\text{such that } x_i(\omega) = \sum_{j \in I} x_{ji}(\omega) \quad (8)$$

$$x_{ji}(\omega) \leq \chi_j M_{ji}(\omega) \quad \forall j, \quad (9)$$

and the production function (5). Here, Y_i is income, and \mathbb{P}_i is the price index in region i . $\chi_j \leq 1, j \in I$ are the shock realizations. We assume the shocks destroy some of the

¹¹Appendix E.2 considers a CES aggregate of inputs.

orders of inputs, M_{ji} , that have been placed in the region in the first stage, and so if a shock materializes, the firm receives fewer inputs than its order. This captures the notion that risk is associated with a disruption of the quantity of inputs that arrive for production for reasons that can include climate-associated shocks such as rainfall or floods, and we will calibrate the shock size to match our event study estimates in Section 2.2. We assume the stochastic shocks are origin-specific, and so they affect orders of inputs from all buying regions. As the shocks are not idiosyncratic, they will potentially affect aggregate outcomes.¹²

Note that as second-stage profits (7) are monotonically increasing in input usage $x_i(\omega)$, the firm will always optimally use all available inputs that are delivered of its orders $M_{ji}(\omega)$. In other words, Equation (9) will always hold with equality.

The first order conditions of the firm's second stage problem (7) pin down a firm's optimal choices of labor l_i , as well as its price p_i , quantity q_i , and profits π_i as a function of the vectors of first stage orders $\mathbf{M}_i = \{M_{ji}\}_{j=1}^I$ and origin-specific shocks, $\boldsymbol{\chi} = \{\chi_j\}_{j=1}^I$. In particular, the expression of profits for a firm in region i , suppressing the variety index ω for concise exposition, is:

$$\pi_i(\mathbf{M}_i; \boldsymbol{\chi}) = \left[\frac{\sigma(1-\beta) + \beta}{\beta(\sigma-1)} \right] \left[\frac{\beta(\sigma-1)}{\sigma} \right]^{\frac{\sigma}{\beta+\sigma(1-\beta)}} w_i^{\frac{\beta(1-\sigma)}{\beta+\sigma(1-\beta)}} \left[\mathbb{Y}_i \mathbb{P}_i^{\sigma-1} \right] \phi_i^{\sigma-1} \left(\sum_{j \in I} \chi_j M_{ji} \right)^{(1-\beta)(\sigma-1)} \left. \right]^{\frac{1}{\beta+\sigma(1-\beta)}}. \quad (10)$$

First stage. In the first stage, prior to the realization of shocks, final goods producers in all locations choose their orders M_{ji} of inputs to maximize expected profits. Firms have rational expectations and make their input sourcing decisions based on the true joint distribution of origin-specific disruption shocks, $G(\boldsymbol{\chi})$.¹³ While the model can readily accommodate alternative belief structures, the assumption of rational expectations is useful for our estimation approach. We consider an alternative belief structure in Section 4.5.

The firm's problem in stage one is

¹²These origin-specific shocks can alternatively be viewed as a disruption to all trade costs/transport routes with the shocked region (Balboni et al., 2024).

¹³In our quantification of the model, we assume that these shocks are binary, occurring with probability ρ_i in each location i , and we permit spatially-correlated shocks.

$$\max_{M_i \geq 0} \mathbb{E}_{\chi} \left[\lambda_i \left(\pi_i(\mathbf{M}_i; \boldsymbol{\chi}) - \sum_{j \in I} p_{ji}^M M_{ji} \right) \right], \quad (11)$$

where p_{ji}^M is the order cost of inputs from j in i , and $\pi_i(\mathbf{M}_i; \boldsymbol{\chi})$ is as in Equation 10. The first order conditions of this problem are

$$\mathbb{E}_{\chi} \left[\lambda_i \left(\chi_j \Theta_i \left[\sum_{j \in I} \chi_j M_{ji} \right]^{\frac{-1}{\beta + \sigma(1-\beta)}} - p_{ji}^M \right) \right] \leq 0 \quad \forall j, \quad (12)$$

where $\Theta_i = \left[\frac{(1-\beta)}{\beta} \right] \left[\frac{\beta(\sigma-1)}{\sigma} \right]^{\frac{\sigma}{\beta + \sigma(1-\beta)}} w_i^{\frac{\beta(1-\sigma)}{\beta + \sigma(1-\beta)}} \left[[Y_i \mathbb{P}_i^{\sigma-1}] \phi_i^{\sigma-1} \right]^{\frac{1}{\beta + \sigma(1-\beta)}}$ is a function of equilibrium aggregates that are potentially stochastic, as Y_i , w_i , and \mathbb{P}_i might depend on the shock realizations across regions.

These first-order conditions highlight that when placing an order for intermediate inputs of a given origin j , firms equate expected marginal benefits and marginal costs. Moreover, this optimality condition elucidates under which circumstances the firm does not source from a particular location. This occurs if the expected marginal benefit from placing an infinitesimal order in location j , with optimal orders elsewhere, is strictly smaller than its expected price, $p_{ji}^M(\boldsymbol{\chi})$.

Proposition 1 *Ex-ante profits are concave in orders of inputs M_{ji} .*

Proof. See Appendix C. ■

This property of the firm's problem, which arises from the firm's inability to adjust input orders ex-post, together with downward-sloping final demand for the firm's good, is important for the firm's optimal sourcing strategy. Interestingly, it implies that the firms behave as if they are risk-averse, even without explicit risk aversion in managerial preferences, when placing their inputs orders to maximize expected profits. As a result, the "risk aversion" from the concavity in profits implies firms will optimally diversify sourcing locations.

Note that our setting does not feature standard love-for-variety motives for diversification. However, while the assumption of perfect substitutability of inputs is stark, inputs from different locations are differentiated by their risk profiles. Appendix E.2 shows that the concavity of firm profits continues to hold with a CES aggregator of inputs from different origins, featuring love-for-variety effects. Our baseline

assumption permits sharp analytical insights and allows us to focus purely on the risk-diversification motive.¹⁴

3.2 General Equilibrium

In the second stage, shocks are realized, inputs are delivered across regions, and goods and labor markets clear. The labor market clearing condition for region i is

$$\underbrace{L_i - \frac{\bar{M}_i}{z_i}}_{\tilde{L}_i, \text{ Final goods labor}} = \left[\frac{\beta(\sigma - 1)}{\sigma} \frac{1}{w_i} [Y_i \mathbb{P}_i^{\sigma-1}]^{\frac{1}{\sigma}} \left(\phi_i \left(\sum_{j \in I} \chi_j M_{ij} \right)^{1-\beta} \right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\beta + \sigma(1-\beta)}}, \quad (13)$$

where \tilde{L}_i is the labor used in the production of final goods in i , and $\frac{\bar{M}_i}{z_i}$ is the labor used in the production of $\bar{M}_i = \sum_{j=1}^J \tau_{ij} M_{ij}$ intermediates to ship to all regions $j \in I$ from region i . Goods markets clear in each region, implying that the region's income is equal to its expenditure:

$$Y_i = w_i L_i + \hat{\Pi}_i, \quad (14)$$

where $\hat{\Pi}_i$ are the aggregate profits in i of the final goods firms as in Equation (7) less their intermediate goods order costs

$$\hat{\Pi}_i = \int \pi_i(\omega) d\omega - \int \sum_j p_{ij}^M M_{ij}(\omega) d\omega. \quad (15)$$

Notice that we assume firms pay for their orders of intermediate inputs, not for the fraction they receive after the shock. Additionally, Equation (13) implies that the full quantity of intermediates ordered in stage 1 is produced. This implies that the shocks “destroy” a fraction of produced inputs.^{15,16} The equilibrium of the economy

¹⁴With a finite elasticity of substitution, firms would choose to source from all locations, inconsistent with the data on sourcing shares, which features many zeros. Here, the diversification motive would imply they source more at the intensive margin from each region. To match the observed sourcing shares, the model would then have to have fixed costs of sourcing, rendering it intractable.

¹⁵We do not observe actual contracts between firms in the data, so we have to make an assumption regarding what fraction of the orders of inputs are paid for. Our setup would remain tractable under alternative assumptions, e.g. only a fraction of the order is paid for upfront. While that would change the input costs entering Equation (11), it would not change the concavity of first stage profits in order costs, which is the key mechanism for firm input diversification in this framework.

¹⁶We assume all inputs ordered are produced, while our event studies showed a decline in the sales of affected firms. Our model is consistent with this pattern as the event studies are based on the

is formally defined in Appendix C.1.

Features of the equilibrium. As all firms in a region are homogeneous, under the unit mass of firms assumption, the regional price index $\mathbb{P}_i = p_i$, and aggregate profits $\hat{\Pi}_i = \hat{\pi}_i$. We can then characterize several features of the equilibrium.

Lemma 1 *Aggregate profits are a constant fraction of labor income $\hat{\Pi}_i = \frac{1}{\sigma-1}w_iL_i$. Further, aggregate expenditure on materials in i is given by*

$$\sum_j p_{ij}^M M_{ij} = (1 - \beta)w_iL_i, \quad (16)$$

and aggregate income in location i is given by,

$$Y_i = \frac{\sigma}{\sigma - 1}w_iL_i. \quad (17)$$

Proof. See Appendix C.2. ■

Lemma 2 *The aggregate labor demand of final goods producers is inelastic, independent of the realization of shocks, χ , and is a constant share of the aggregate labor endowment,*

$$\tilde{L}_i = \beta L_i. \quad (18)$$

Proof. See Appendix C.2. ■

To understand the intuition behind Lemma 2, consider the case of firms in a region facing negative shocks in its sourcing locations at the start of stage 2. Due to input disruptions, all else equal, the demand of final goods producers for labor falls. But in equilibrium, this decline is exactly offset by the increase in final goods prices and real wage declines, as aggregate consumer demand is downward-sloping. The net effect is that the aggregate labor demand from final goods producers remains unaffected.

Equation (18) shows that equilibrium wages must be such that the remaining workers are used by the intermediate inputs sector in stage 1. This implies that equilibrium wages w_i and input prices p_i^M are such that stage 1 firm input orders demand $(1 - \beta)L_i$ to produce \bar{M}_i so that labor markets clear.

value of inputs that are shipped after a shock. They do not speak to the quantity of inputs produced or the payments firms have made for their input orders.

Lemma 3 *Let labor in region 1 be the numeraire. Equilibrium relative wages w_i are deterministic. This implies that aggregate income in location i is also deterministic.*

Proof. See Appendix C.2. ■

Lemma 3 follows immediately from the discussion above. There is a unique wage w_i in each location such that equilibrium input orders placed by firms in stage 1 require $(1 - \beta)L_i$, the labor not used in final goods production in stage 2, to be produced. This result is not imposed by assuming wages are predetermined or fixed in stage 1. Rather, since aggregate labor demand from final goods producers is perfectly inelastic and independent of realized shocks, the equilibrium vector of regional wages, w_i , must be determined entirely by conditions prevailing before uncertainty is resolved. Consequently, there exists exactly one deterministic wage vector that clears regional labor markets, equating the labor demand of intermediate goods producers with the labor supply net of the invariant labor requirements for final goods production.

This simplifies the analysis substantially: while wages could potentially vary across states of the world, by Lemmas 1-3, wages, input prices, nominal income and profits are deterministic. The only aggregate variable that is stochastic, varying with the realization of shocks, is the ideal price index, \mathbb{P}_i

In the ex-post general equilibrium, the expression for Θ_i , which is part of the marginal contribution to profits of a marginal unit of M_{ji} (Equation 7), is given by the following expression:

$$\Theta_i = (1 - \beta)w_iL_i \left(\sum_{j \in I} \chi_j M_{ji} \right)^{-\frac{(1-\beta)(\sigma-1)}{\beta+\sigma(1-\beta)}}.$$

This implies that Θ_i is stochastic from the perspective of firms in stage 1.

Ex-Ante General Equilibrium As pointed out above, the vector of relative wages is deterministic and determined at the first stage, and intermediate goods producers employ $(1 - \beta)L_i$ workers in input production. In turn, due to the linear technology assumption, it must be the case that in equilibrium, the production of intermediates in each location is equal to $\bar{M}_i = (1 - \beta)z_iL_i$. In the equilibrium of this economy, the vector of wages, $\{w_i\}_{i=1}^I$, must be such that total demand from intermediate goods producers in each region exactly equals this amount.

From trade balance and optimal total intermediate-expenditure conditions, we

derive the following equilibrium system, generating the equilibrium wage vector:

$$w_j L_j = \sum_i w_i L_i s_{ji}(\{w_i\}_{i=1}^I) \quad ; \quad s_{ji}(\{w_i\}_{i=1}^I) = \frac{\frac{w_j \tau_{ji}}{z_j} M_{ji}(\{w_i\}_{i=1}^I)}{\sum_\ell \frac{w_\ell \tau_{\ell i}}{z_j} M_{\ell i}(\{w_i\}_{i=1}^I)} \quad \forall j \in I,$$

where crucially, the matrix of sourcing shares defined by $\left\{s_{ji}(\{w_i\}_{i=1}^I)\right\}_{i=1, j=1}^I$ is a function of the vector of wages, the parameters of the model, and the probability distribution of the shocks.¹⁷ This completes the description of the economy.

Welfare Agents' welfare is given by expected consumption, which is equal to the final goods producers' output and varies by region. In general equilibrium, the aggregate output of the final sector in region i , conditional on available inputs, is

$$Q_i(\mathbf{M}_i; \boldsymbol{\chi}) = \phi_i \beta^\beta L_i^\beta \left(\sum_j \chi_j M_{ji} \right)^{1-\beta},$$

and expected welfare becomes

$$\mathcal{W}_i = \mathbb{E}_{\boldsymbol{\chi}} [\log Q_i(\mathbf{M}_i; \boldsymbol{\chi})] = \log \phi_i + \beta \log \beta + \beta \log L_i + \mathbb{E}_{\boldsymbol{\chi}} \left[(1 - \beta) \log \left(\sum_j \chi_j M_{ji} \right) \right]. \quad (19)$$

As is clear from this welfare expression, since consumers are risk averse under log utility, the sourcing strategy selected by the final goods producers has effects on their welfare. Consumers benefit from diversification in firms' sourcing strategies.¹⁸

3.3 A Two Location Example

To gain intuition, consider a simple case with two locations. Region 1 is risky and receives a shock $\chi_1 < 1$ with probability ρ , and region 2 is a safe location.¹⁹ Additionally, there are no trade costs, and therefore, the optimal intermediate bundle chosen by firms is the same in both locations.

¹⁷Similar non-linear systems of equations in wages appear in several static trade models. Note that here, the system includes orders of intermediates, M_{ji} , which are also equilibrium objects and do not have a closed-form solution.

¹⁸The model can be solved under CRRA preferences, which can be parameterized to imply stronger risk aversion and larger welfare gains from diversification.

¹⁹That is, $\mathbb{E}_{\boldsymbol{\chi}}^1 = \rho \chi_1 + (1 - \rho)$ and $\mathbb{E}_{\boldsymbol{\chi}}^2 = 1$.

Notice that in equilibrium it must be that $p_1^M < p_2^M$, because otherwise, the safe location's input is unambiguously better than the input from the risky location, and the labor market will not clear in the risky location.²⁰

The optimal stage 1 sourcing choices for firms from both regions $i \in \{1, 2\}$ is

$$M_{1i} : \rho \chi_1 \lambda_i^S \Theta_i^S [\chi_1 M_{1i} + M_{2i}]^{\frac{-1}{\beta + \sigma(1-\beta)}} + (1 - \rho) \lambda_i^{NS} \Theta_i^{NS} [M_{1i} + M_{2i}]^{\frac{-1}{\beta + \sigma(1-\beta)}} = p_1^M \quad (20)$$

$$M_{2i} : \rho \lambda_i^S \Theta_i^S [\chi_1 M_{1i} + M_{2i}]^{\frac{-1}{\beta + \sigma(1-\beta)}} + (1 - \rho) \lambda_i^{NS} \Theta_i^{NS} [M_{1i} + M_{2i}]^{\frac{-1}{\beta + \sigma(1-\beta)}} = p_2^M, \quad (21)$$

where $\Theta_i^S = \frac{(1-\beta)(\sigma-1)}{\sigma} Y_i (\chi_1 M_{1i} + M_{2i})^{-\frac{(1-\beta)(\sigma-1)}{\beta + \sigma(1-\beta)}}$ and $\Theta_i^{NS} = \frac{(1-\beta)(\sigma-1)}{\sigma} Y_i (M_{1i} + M_{2i})^{-\frac{(1-\beta)(\sigma-1)}{\beta + \sigma(1-\beta)}}$. As discussed above, Θ_i is stochastic, and depends on whether or not the shock materializes in region 1. Under the monopolistic competition assumption, all firms take these aggregates as given. Entering these shifters into the first order conditions of the firms, we can solve for optimal orders as a function of wages:

$$M_{1i} = \frac{(1-\beta)(\sigma-1)}{\sigma} Y_i \left[\frac{1-\rho}{p_1^M - \chi_1 p_2^M} - \frac{\rho}{p_2^M - p_1^M} \right] \quad (22)$$

$$M_{2i} = \frac{(1-\beta)(\sigma-1)}{\sigma} Y_i \left[\frac{\rho}{p_2^M - p_1^M} - \frac{(1-\rho)\chi_1}{p_1^M - \chi_1 p_2^M} \right]. \quad (23)$$

Let wages in the less-risky region 2 be the numeraire. As intermediates are priced at marginal cost and from the labor market clearing condition (Equation 13), a constant fraction of labor is used in the production of intermediates, and we can show that equilibrium wages in the risky region 1 are given by

$$w_1 = \frac{z_1 z_1 L_1 \chi_1 + z_2 L_2 (1 - \rho(1 - \chi_1))}{z_2 z_1 L_1 (\rho + \chi_1 (1 - \rho)) + z_2 L_2}. \quad (24)$$

Equation 24 shows that the nominal wage in the risky location relative to the safe one is a function of relative productivities, relative sizes, and the probability and magnitude of the shock. This wage is increasing in relative productivity and decreasing in relative population of location 1, and particularly relevant to our application, decreasing in both the probability and the magnitude of the sourcing disruption.

²⁰The fact that in this simple case, we have an interior solution for firms in both locations does not need to hold in general when there are multiple locations and trade costs.

3.4 Comparative Statics

For a larger number of regions, the model does not have an analytical solution, so we first illustrate the model’s properties in a stylized 3-region setting. We assume the regions are homogeneous in firm productivity ϕ_i , labor endowment L_i , and intermediate producer productivity z_i . Trade is costly between regions with distance elasticity of 0.5. We assume if a shock occurs, 90% of the inputs are destroyed ($\chi = 0.1$).

To focus on spatial variation in risk, we assume the three locations are equidistant, but risk varies across space. We assume $\frac{1}{I} \sum_{i=1}^3 \rho_i = 0.5$ and contrast costly trade to autarky. Appendix C.3 allows for regions to vary in their distance to each other, placing them on a straight line, but with constant risk ($\rho_i = 0.5$ for all i).

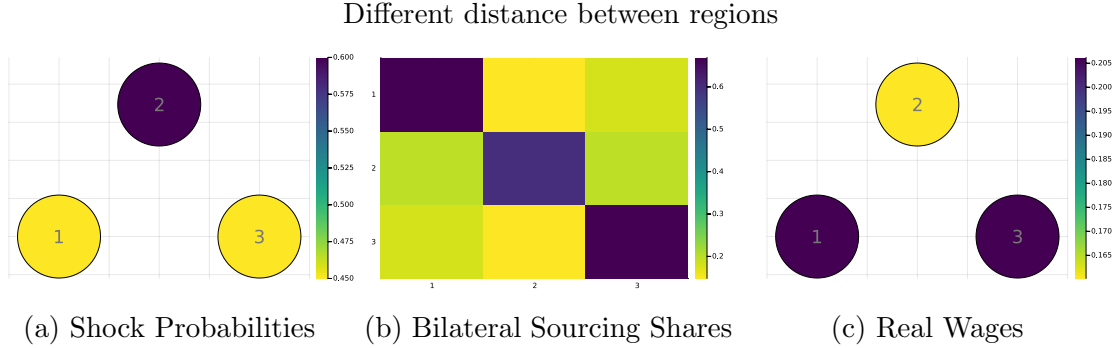
Heterogeneous risk, homogeneous distance. The left panel of Figure 3 illustrates the regional maps and the shock probabilities of each region in the heterogeneous risk case. As regions are equidistant, geography does not play a role in diversification. The middle panel shows the bilateral sourcing shares between regions. The diagonal is the darkest: in the presence of trade costs, all regions source most of their inputs from their own region despite heterogeneous risk. However, there is clear variation. Regions 1 and 3 (the safest regions), see the most “own sourcing.” The riskiest region 2 diversifies the most. All regions source inputs from other regions, with relatively larger shares from those with low risk.

The right panel shows that expected real wages across regions are negatively correlated with shock risk, and are highest in safest locations despite identical regional fundamentals. The underlying mechanisms are that safer regions experience higher labor demand for their intermediate inputs from all regions, pushing up nominal wages. They also face a lower price index of their final goods, as they can source safer “domestic” inputs without paying trade costs. Notice that in general equilibrium, the wage impacts on riskier regions will modulate sourcing from them.

Heterogeneous risk and autarky. In the same environment, we set trade costs to infinity, shutting down inter-regional input sourcing. Appendix Figure C1 illustrates that in regional autarky, the riskiest region sees the lowest expected real wages, while the safest regions see the highest expected real wages. These regions have the lowest expected prices due to lower shock probabilities and fully domestic sourcing.

We next consider how expected real wages change across regions moving from costly trade to autarky in Panel A, Figure 4. Interestingly, all regions see a decline

Figure 3: Scenario with heterogeneous risk, homogeneous distance



Note. The figures in the left panel show the probability that each region is hit by a shock, as well as a visual representation of the geographical location of regions in space. The figures in the middle panel consist of a 3x3 input-output matrix where the buying regions are on the vertical axis, and the supplying regions are on the horizontal axis. Each line represents the share of inputs purchased by a buying region from each supplying region. The right panel presents the real wages for each region. Regions are equidistant from each other. The scales are shown to the right of each figure.

in expected real wages moving to trade from autarky. In this setting, there are no gains from varieties. The primary reason for trade is for risk diversification. However, trade is costly, so the benefits of diversification are obtained at a higher average input price, raising regional price indices and lowering expected real wages.

The lower expected real wages under costly trade do not imply welfare losses from trade: Panel B of the figure illustrates that there is a large decline in the volatility of real wages under trade for all regions. Supply chain diversification lowers the variance in final goods prices across all regions, insuring against shocks and real wage volatility. Household welfare is the *expected* log quantity of the final goods bundle consumed (Equation 19). As a result, the decline in variance of real wages contributes positively to their welfare, offsetting the decline in expected real wages, and trade is welfare-improving.²¹

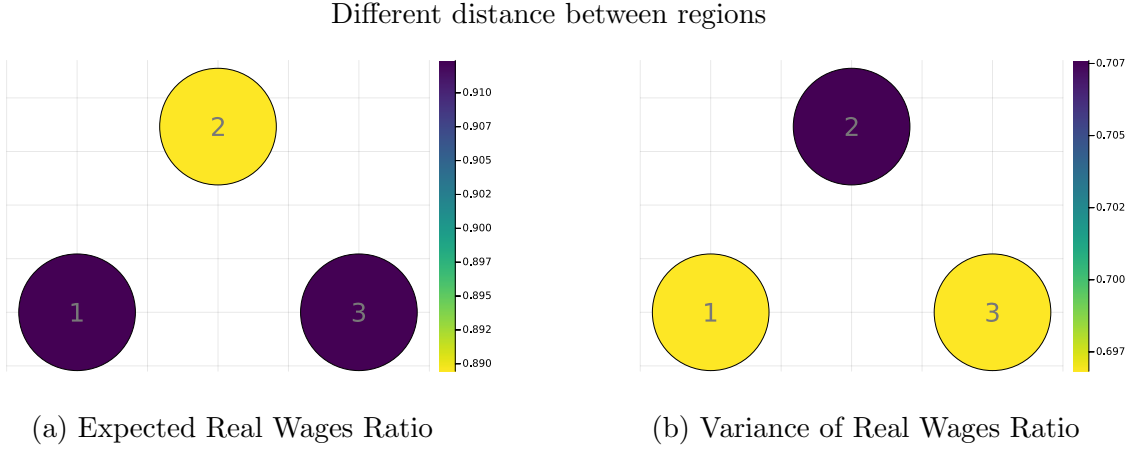
4 Quantification

4.1 Solution Approach

The solution to the quantitative model introduced in Section 3 requires overcoming three computational challenges. First, the perfect substitutability across intermediate inputs from different origins, combined with the existence of trade costs, implies that

²¹Recall, $\mathbb{E}[\log X] \approx \log \mathbb{E}[X] - c\mathbb{V}[X]$. This result depends on the assumption of log utility.

Figure 4: Comparison between heterogeneous risk under costly trade and autarky



Note. In this figure we plot the expected real wages (left panel) and variance of real wages (right panel) for the scenario with heterogeneous risk and costly trade shown in Figure 3 relative to the scenario with heterogeneous risk and autarky shown in Figure C1. The variance of real wages is computed across potential states of the world. Here, regions are equidistant from each other. The scales are shown to the right of each figure.

the solution to the firms' sourcing problem may not necessarily be interior; that is, firms in some regions might find it optimal not to source from certain origins. Second, finding the solution to the firms' optimal sourcing problem involves computing a high-dimensional expectation over 2^I states of the world.²² Third, the two challenges mentioned above are compounded by the need to find the equilibrium of the model, which amounts to finding the vector of wages for which all markets clear.

Given a vector of wages, $\{w_i\}_{i=1}^I$, and shock probabilities, $\{\rho_i\}_{i=1}^I$, we leverage the structure of the model to solve it efficiently. The first property of the problem described in Equation 11 is that the objective function is concave, and that the constraints are linear. Thus, any locally optimal point is also globally optimal, i.e., the Karush-Kuhn-Tucker (KKT) conditions are both necessary and sufficient for global optimality. These allow us to solve the firm's problem by combining the stationarity and complementary slackness conditions to find that at the optimum, the following condition holds with equality:

²²There are more than 600 districts in India, but for computational feasibility we group small contiguous districts to create 271 regions. We implement our model for the 271 regions, so that involves computing expectations over $2^{271} \approx 10^{82}$ states of the world.

$$\mathbb{E} \left(\chi_i \Theta_j \left[\sum_{i=1}^I \chi_i M_{ij} \right]^{\frac{-1}{\beta + \sigma(1-\beta)}} \right) M_{ij} = \frac{w_i \tau_{ij}}{z_i} M_{ij} \quad \forall i \in I,$$

which results from multiplying the first order condition in Equation 12 by M_{ij} . We substitute for the general equilibrium object, Θ_j , to derive the following:

$$(1 - \beta) w_j L_j M_{ij} \mathbb{E} \left[\chi_i \left(\sum_{i=1}^I \chi_i M_{ij} \right)^{-1} \right] = \frac{w_i \tau_{ij}}{z_i} M_{ij} \quad \forall i \in I.$$

This system of I equations in I unknowns defines a nonlinear complementarity problem for which efficient numerical optimization routines exist.²³ Finally, we approximate the high-dimensional expectation by using simulations, effectively solving the following system of equations for each region:²⁴

$$(1 - \beta) w_j L_j M_{ij} \frac{1}{S} \sum_{s=1}^S \left[\chi_i^{(s)} \left(\sum_{i=1}^I \chi_i^{(s)} M_{ij} \right)^{-1} \right] = \frac{w_i \tau_{ij}}{z_i} M_{ij} \quad \forall i \in I.$$

The procedure described above yields a solution to the firms' sourcing problem given a vector of wages, $\{w_i\}_{i=1}^I$. To find the equilibrium wages, we manipulate the trade balance and the optimal total intermediates expenditure conditions to derive the following equilibrium system,

$$w_j L_j = \sum_i w_i L_i s_{ji}(\{w_i\}_{i=1}^I) \quad ; \quad s_{ji}(\{w_i\}_{i=1}^I) = \frac{\frac{w_j \tau_{ji}}{z_j} M_{ji}(\{w_i\}_{i=1}^I)}{\sum_k \frac{w_k \tau_{ki}}{z_k} M_{ki}(\{w_i\}_{i=1}^I)} \quad \forall j \in I,$$

where, the matrix of sourcing shares defined by $\left\{ s_{ji}(\{w_i\}_{i=1}^I) \right\}_{i=1, j=1}^I$ is a function of the vector of wages and model parameters. The solution to the system of equilibrium conditions above finds the equilibrium wages conditional on a vector of probabilities, $\{\rho_i\}_{i=1}^I$. We describe how we calibrate these probabilities in the next subsection.

²³We solve this problem using the optimizer `PATH` implemented on `Julia` through the optimization modeling language `JuMP`.

²⁴In our estimation procedure and in the computation of counterfactuals, we use 10000 simulations.

4.2 Calibration

For computational feasibility, we group the over 600 districts in India into 271 regions by grouping contiguous low-population districts.²⁵ We calibrate our model to these 271 regions. To calibrate the model for India as a whole, we complement our transaction data with the Annual Survey of Industries (ASI), which is a nationally representative survey of manufacturing plants in India with more than ten employees. We primarily use the wave of 2006-7 since it is the last year for which the ASI has publicly available district identifiers, and more recent years cannot be used at the district-level to calibrate a spatial model.

We need to calibrate the following parameters and moments: the demand elasticity (σ), labor endowments by district (L_i), regional productivities (ϕ_i, z_i), the labor share in the production function (β), iceberg trade costs (τ_{ij}), the input disruption due to the shock (χ_j) and flood probabilities (ρ_i).²⁶

First, we set the demand elasticity $\sigma = 2$ following [Boehm et al. \(2023\)](#). Second, we use the ASI to obtain employment by district, which is our labor endowment, L_i . We use the documented evidence in Fact 4 in Section 2.2 to choose the input disruption parameter χ_j , and match the drop in buyer purchases from the event study. This generates a response to the incidence of a disruption within our model that matches the drop estimated in Fact 4.²⁷

Productivities To estimate productivities by district, ϕ_i , and the labor share β , we follow the production function estimation literature and use the [Akerberg, Caves, and Frazer \(2015\)](#) approach (henceforth ACF).²⁸ We use revenues as the dependent variable and labor, materials, and capital as production function inputs and estimate

²⁵We aggregate districts with fewer than 10000 manufacturing workers to a single district within a state, or merge them to neighboring larger districts in their own state.

²⁶Calibrating these parameters prevents us from using geographical units that are smaller than districts, as additional data for calibration are not available for smaller areas.

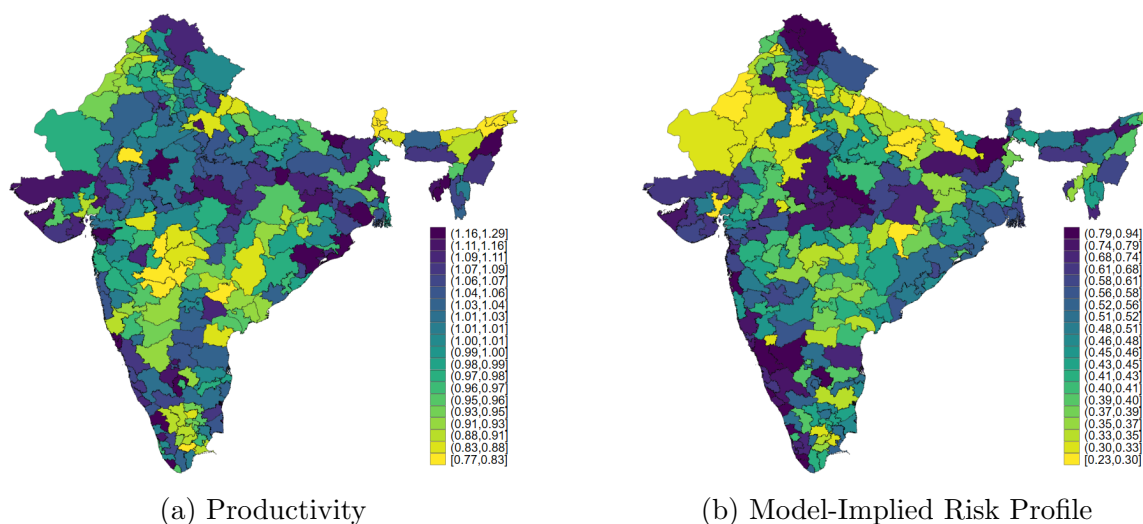
²⁷Note that we calibrate χ_j to the impact of the incidence of a flood. The disruption probabilities in our model capture many sources of risk, climate- and non-climate-related, as we discuss below. This calibration assumes all disruptions, if they occur, are as severe as the realization of flood events. We do not have other exogenous shocks to discipline the severity of other sources of risk, but we can readily assess robustness to alternative values of χ_j in the quantification.

²⁸This approach requires lagged values of labor and materials as instruments, and we need a panel of firms. However, the public version of the ASI is a cross-section of plants which prevents constructing a firm-level panel. As a solution, we use the waves for 2004-05, 2005-06, and 2006-07 to construct a synthetic panel at the industry-district level. We then treat each industry-district pair as a “firm” for the purposes of estimation.

the production function parameters and the productivities.²⁹

Panel A of Figure 5 illustrates the estimated variation in district-level productivities. From the ACF procedure, we also get the corresponding coefficients for labor, materials, and capital. The results are shown in the left panel of Table 3, where the materials share is 0.81, the labor share is 0.17, and the capital share is 0.08. We compute the labor share as $\beta = 1 - 0.81 = 0.19$. As we do not have capital in the model, we think of the labor share as the share of capital-augmented labor, so we include both capital and wage expenses into the calculations.

Figure 5: Estimated productivities and disruption probabilities



Note. In this figure, we plot the estimated district-level productivities (left panel) and the model-implied district-level disruption probabilities (central panel). Productivities are estimated using the ACF procedure as described in the text. Baseline disruption probabilities are obtained by matching model-implied sourcing shares to the data as described in the text. The right panel plots the district-level disruption probabilities implied by the parameterized approach outlined in the text. The scales are shown to the right of each figure.

Iceberg trade costs The iceberg trade costs τ_{ij} are estimated using our transaction data, leveraging our information on transaction-level prices. Our data is only available if one node of the transaction lies in one particular state, but we need to back out trade costs for each bilateral pair of districts throughout India. To address this, we proceed in two steps. First, we use our transaction data, focus on firms in our state

²⁹Once we back out the ACF productivity for each industry-district pair, we aggregate at the region level by using weights based on the relative importance of each industry in each region. In the few cases where productivity cannot be estimated due to missing data for smaller districts, we assign those regions the average productivity of their closest neighbors.

that sell their goods, and aggregate the data at the seller-buyer-product-time level. We then estimate Equation 25.

$$\log(p_{s,b,t,q}) = \gamma_1 \log(\text{distance}_{s,b}) + \gamma_2 \mathbb{1}(b \text{ in same state as } s)_{s,b} + \delta_{s,q,t} + \epsilon_{s,b,t,q}, \quad (25)$$

where $p_{s,b,t,q}$ is the price charged by seller s to buyer b for product q at time t . For each buyer-supplier pair, we compute the log distance between them as reported in our transaction data. We also include an indicator variable for whether the buyer (b) is in our state. The coefficient on distance captures how prices charged change as distance increases. Importantly, we add seller-product-time fixed effects $\delta_{s,q,t}$, so effectively, the coefficients γ_1 and γ_2 are being identified by sellers that sell the same product to multiple buyers in a given time period.

The underlying assumption is that unobserved iceberg trade costs τ_{ij} are proportional to distance. In our data, the same seller charges different prices to different buyers for the same product and month. We assume this variation partly depends on unobserved iceberg costs. As we include seller-product-time fixed effects, the estimates are not driven by seller-shocks (e.g., productivity) that may affect prices.

Note that iceberg trade costs conventionally include observed costs such as freight or transportation, but additionally other unobserved costs such as contracting frictions, linguistic/ethnic differences, unobserved preference shifters, etc. We assume these are proportional to observed distance, which is common in gravity estimation. In our data, freight costs are not required to be included in the values of goods shipped reported, though sellers might include this. So, it is likely our iceberg trade cost estimation might not include transportation costs. Some sellers also explicitly separately report freight costs. As robustness, we also create an ‘‘Adjusted Price’’ measure, which adds the reported freight costs. Estimates remain similar, but the sample is much smaller as this variable is missing for many observations in the data. The results of this regression can be found in the right panel of Table 3.

We then use the estimated coefficient to predict trade costs for the rest of India. We compute bilateral distances between district centroids and predict trade costs between regions using the estimated coefficients $\hat{\gamma}_1$ and $\hat{\gamma}_2$. We assume that the border effect estimated through coefficient $\hat{\gamma}_2$ is the same for all states.

Disruption probabilities. Our model implies that bilateral sourcing shares are pinned down by district fundamentals like productivities and labor force and by bi-

Table 3: Estimation results

Panel A: Production Function Estimation		Panel B: Trade Costs Estimation		
	log(Sales)		log(Price _{s,b,t,q})	log(Adj. Price _{s,b,t,q})
log(Materials)	0.81*** (0.076)	log(distance from s to b)	0.0174*** (0.0001)	0.0186*** (0.0002)
log(Workers)	0.17*** (0.061)	$\mathbb{1}(b \text{ in same state as } s)$	-0.086*** (0.0001)	-0.0798*** (0.0009)
log(Fixed Capital)	0.08 (0.063)			
Number of Observations	9128	Number of Observations	65,477,898	45,338,641

Note. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$ Panel A presents the results of the production function estimation using the ACF procedure. The reported coefficients are for log materials, log number of workers, and log fixed capital as calculated from the ASI. Panel B presents the results for the trade costs estimation using our transaction data. The outcome is the log price charged by a seller in our state (s), for a given product (q), to a buyer (b) in a given month-year period (t). The main regressors are log distance from buyer to seller and a dummy that takes the value of 1 if the buyer is in the same state as the seller. We control for seller-product-time fixed effects. In column 2 of Panel B, we compute the adjusted price by adding the total transaction value and “other” reported costs (including freight), and dividing by quantity. Other costs include additional self-reported transportation costs not reported in the transaction value.

lateral trade costs, as well as the vector of district-level shock probabilities. Therefore, we can obtain the vector of shock probabilities, ρ_i , by minimizing the distance between the observed sourcing shares in the data with those implied by the model. When estimating the probabilities, we allow for spatial correlation in the realization of disruptions, as floods or other disruptions might affect more than one district.³⁰

The intuition of the exercise is as follows: conditional on the rest of the parameters and moments of the model, we pick the shock probabilities of each district to minimize the distance between the model-implied shares with the observed shares of purchases from every district in our state to each other district in India.³¹ This is our baseline

³⁰We assume that these disruptions are generated by a binary random variable that is equal to 1 whenever a normal latent variable with mean 0 and standard deviation 1 is below a threshold equal to $\Phi^{-1}(\rho_i)$, where Φ^{-1} is the standard normal inverse CDF. We allow these latent variables to be correlated across regions, where the correlation in the realizations between region i and region j is equal to $e^{-\zeta \text{Dist}_{ij}}$, where ζ is a measure of spatial decay in this correlation. We estimate ζ in the same routine as the probabilities, ρ_i .

³¹We do not observe the realizations of disruptions in each district, and we remain agnostic on the sources of risk that generate disruptions. However, observed sourcing shares in the data include any realizations of disruptions, which we treat as structural errors. Precisely, given a sourcing strategy in each region, we generate a large number of shocks, χ_i , from the true distribution $\mathcal{G}(\boldsymbol{\chi})$, $\mathcal{P}(\chi_i = \chi) = \rho_i$ and compute the model-implied shock-inclusive sourcing shares. We estimate ρ_i by minimizing the gap between the shares in the data and the average across model simulations, allowing for the spatial correlation as discussed above. Formally, $\min_{\rho \in [0,1]^I} \sum_{j \in I^o} \left(s_{ji}^{\text{Data}} - \frac{1}{S} \sum_{s=1}^S s_{ji} \left(\{\chi_i^s\}_{i=1}^I, \rho_i \right) \right)^2$,

approach, as it allows us to remain agnostic on the sources of risk in the model. Instead, we can validate our model by projecting the shock probabilities on plausible sources of risk. Panel B of Figure 5 plots the baseline disruption probabilities by district. As an alternative approach, we also parameterize regional risk as a function of observables, and estimate the parameters of this function, as described below.

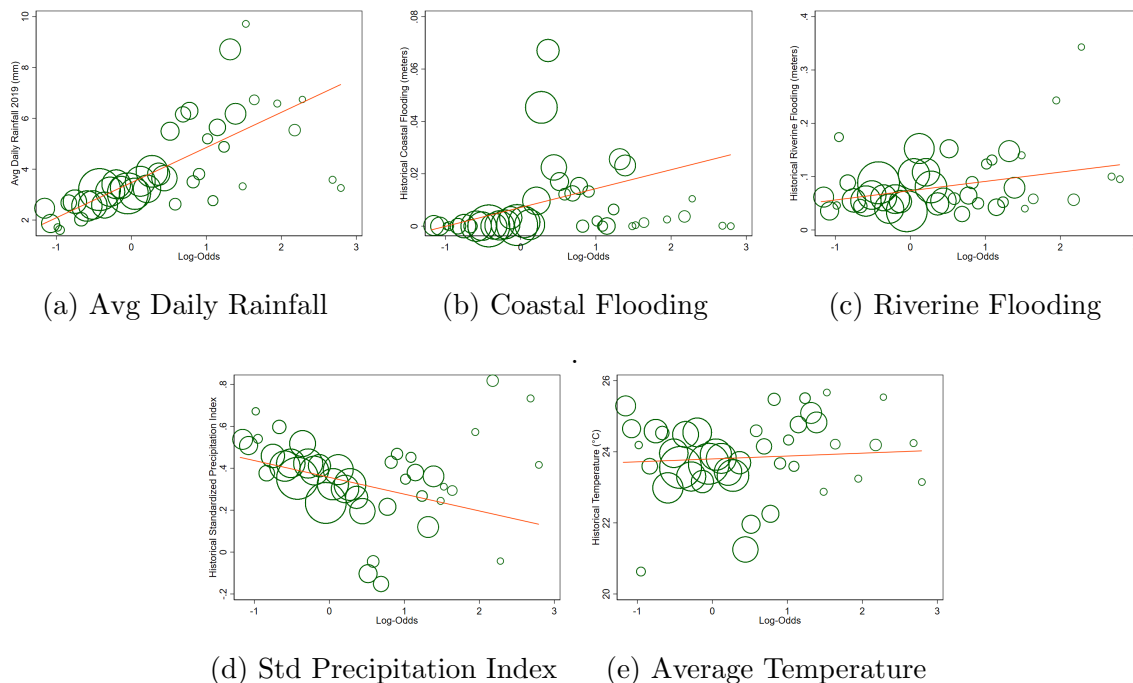
The underlying assumption of our baseline approach is that anything that is not captured by the district-level productivities and trade costs is part of the risk of the district. Of course, in practice, such residuals do not only include flooding risk, but also many other risk components including institutional risk. As these residuals are obtained through a procedure similar to the model-inversion common in trade models, they will also naturally contain model mis-specification, and in particular, other motives for diversification such as love-for-variety. However, in Figure 6 and Appendix Figure D2, we show that our estimated probabilities are significantly correlated with historical and projected average rainfall, coastal flooding, riverine flooding, average temperature, and dryness.

In Table 4, we run regressions of the model probabilities on the climate variables (historical and projected 2050) as well as other variables that could also be related to risk. The climate variables such as daily rainfall, coastal flooding, and average temperature are all strongly significantly correlated with the probabilities, and the R^2 of the regressions are high, around 0.32. To capture institutional features that might affect risk, we add state fixed effects and a district court congestion control in columns 3 and 4. The coefficients on the climate variables remain similar in magnitude and significance. It is also well-known that in a cross-section, more productive regions have lower climate risk. Therefore, significant coefficients on climate variables might simply be picking up the confounding regional productivity effect. While our residuals are estimated conditional on regional productivity, to avoid such confounding, in columns 5 and 6 we additionally include productivity controls such as district productivity and nightlights. The results are similar. Figure D1 further shows that these probabilities do not show a strong correlation with either the estimated district productivities, nor the average distance to the state of our study.

Notice this exercise requires solving jointly for the vector of district-level risk that minimizes the gap between model-implied sourcing shares and data, as all bilateral sourcing shares are equilibrium objects that depend on the fundamentals and risk of

where s is a model simulation with shocks $\{\chi_i^s\}_{i=1}^I$.

Figure 6: Model Probabilities and Historical Observables



Note. We plot the estimated probabilities against historical climate observables. In Figure 6a, we correlate the probabilities with the average daily rainfall in 2019. Figures 6b and 6c use historical coastal and riverine flooding respectively. Figure 6d correlates the probabilities with the standardized precipitation index, a measure of dryness. In Figure 6e, we correlate the probabilities with average temperature. A more detailed definition of each of the variables can be found in Appendix D.1.

other districts. Further, we cannot exactly match all bilateral sourcing shares in the data; we choose a single shock probability for each district, but we observe multiple sourcing shares for that district from all districts in our state. We therefore set up a minimum distance estimator which aims to match the average sourcing shares for each origin district observed across all destination districts in our data. In practice, we match all the bilateral sourcing shares in the data well (Figure 7). As external validation, the right panel of Figure 7 shows that our model also matches the data on sales shares well, which are untargeted moments.³²

³²While our estimated probabilities might seem high, as discussed above, they capture several sources of risk. Further, available evidence from Indian businesses suggests that supply chain disruptions are a key concern. For instance, PwC’s 26th Annual Global CEO Survey in late 2022 found that 50% of India CEOs were concerned about supply chain disruptions (<https://www.pwc.in/assets/pdfs/research-insights-hub/immersive-outlook-3/paradigm-shift-in-supply-chain-management.pdf>).

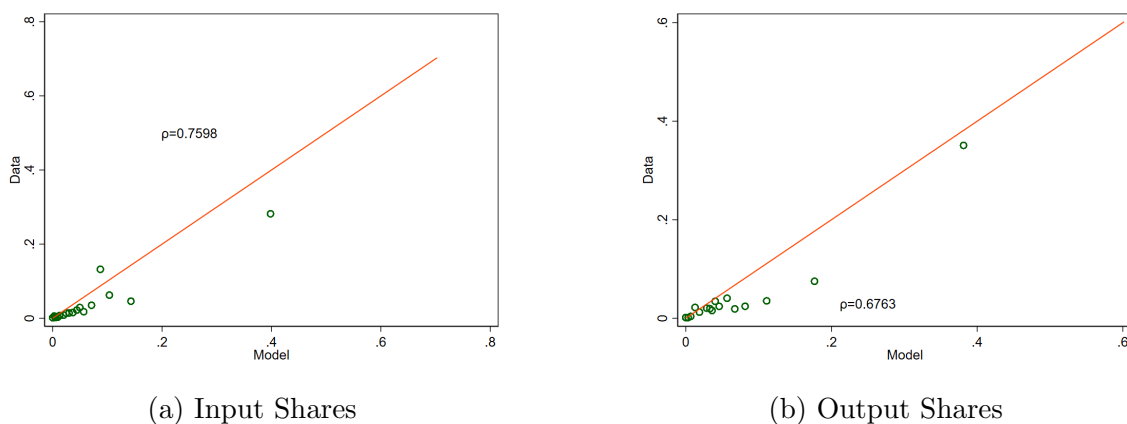
Table 4: Regression of model probabilities on observables

	Historical	Projected (2050)	Historical	Projected (2050)	Historical	Projected (2050)
Daily Rainfall	0.104*** (0.0264)	0.0566*** (0.0182)	0.0981*** (0.0215)	0.0440*** (0.0137)	0.0321 (0.0276)	0.0212 (0.0139)
Coastal Flooding	1.455*** (0.543)	1.418*** (0.311)	2.126*** (0.665)	1.824*** (0.390)	3.066*** (0.576)	1.956*** (0.337)
Riverine Flooding	0.287 (0.337)	0.359 (0.295)	0.216 (0.341)	0.468* (0.276)	0.471 (0.334)	0.591 (0.362)
Avg SPI	-0.155 (0.182)	-0.0444 (0.114)	-0.0936 (0.158)	0.000400 (0.0969)	-0.350** (0.170)	0.0552 (0.177)
Avg Temperature	0.0519*** (0.0175)	0.0669*** (0.0188)	0.0595*** (0.0177)	0.0700*** (0.0181)	0.0852*** (0.0322)	0.0712** (0.0310)
Terrain Controls	Y	Y	Y	Y	Y	Y
Institutional Controls	N	N	Y	Y	Y	Y
Productivity Controls	N	N	Y	Y	Y	Y
State Fixed Effects	N	N	N	N	Y	Y
N	271	271	271	271	271	271
adj. R-sq	0.322	0.313	0.339	0.323	0.367	0.356

Note. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. We estimate regressions of the inverse logit of the estimated model probabilities on observables. In columns 1, 3 and 5 climate variables used are measured with their historical values. In columns 2, 4 and 6 climate variables used are measured with the projected values for 2050. Terrain controls include average elevation and ruggedness, institutional controls include mean court congestion and productivity controls are average nighttime luminosity and our measure of local TFP. Observables are in levels. A more detailed definition of each of the variables can be found in Appendix D.1.

Robustness Our baseline approach has the benefit of remaining agnostic about the sources of disruptions firms face. However, it requires estimating a disruption probability for each district, which is a large number of parameters. As an alternative, we assume that the disruption risk in each district is a function of observables, including the climate and alternative variables in Table 4. We then estimate the coefficients of this function to minimize the distance between model-implied and observed sourcing shares. This has the advantage that we restrict the number of parameters to be estimated to 11. However, as we do not observe all sources of risk, there will be more unexplained variation. Panel A of Figure E6 illustrates district-level disruption risk implied by this approach. Unsurprisingly, as the observable risk measures were correlated with the “agnostic” risk from our baseline approach, the results of the parameterized approach are also correlated with our baseline. Appendix E.1 outlines this approach in more detail, and presents all our quantitative results under this al-

Figure 7: Sourcing shares: Model vs. Data



Note. In this figure, we plot the sourcing shares in the data against the model. The red line is a 45-degree line. In the left panel we plot the input sourcing shares. We target average sourcing probabilities from our state’s districts to the rest of the districts, but we do not force anything to match the particular sourcing shares of each district. The left panel plots each individual district’s input shares. The right panel shows sales shares, which are entirely untargeted. The R^2 of the left panel without the outlier is 0.50 and of the right panel is 0.79 without its outlier.

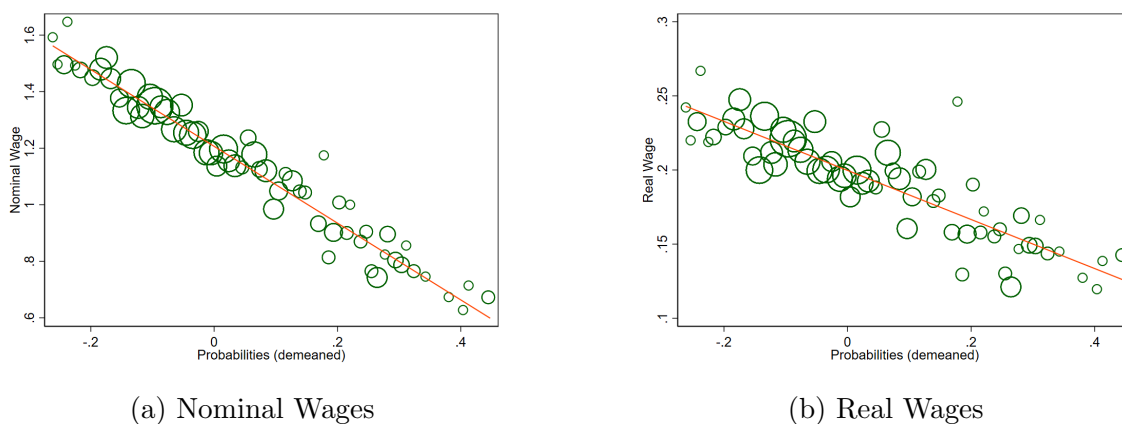
ternative approach. Our main conclusions remain unchanged. Appendix Table D1 summarizes our model calibration.

4.3 Quantitative Results

We first show that the model delivers a strong negative relationship between shock probabilities and relative nominal wages (and real wages) in the cross-section. Figure 8 shows that both nominal and real wages are negatively correlated with shock probabilities, as we would expect. These results quantitatively validate the key trade-off in the model between sourcing risk and input costs and illustrate the baseline distributional consequences of risk: higher-risk regions are poorer in real terms. In Figure D3, we also show that the price index and the variance in real wages are negatively correlated with the shock probabilities.

Probabilities and sourcing shares. To illustrate the rich heterogeneity in bilateral sourcing patterns and disruption probabilities in the quantitative model, we focus on one district, Kolkata, in Figure 9. The left panel illustrates the spatial correlation of disruption probabilities between Kolkata and other districts. The right panel shows the sourcing shares of Kolkata from other districts. Firms diversify, but sourcing strategies depend on geography – they source more from relatively geographically

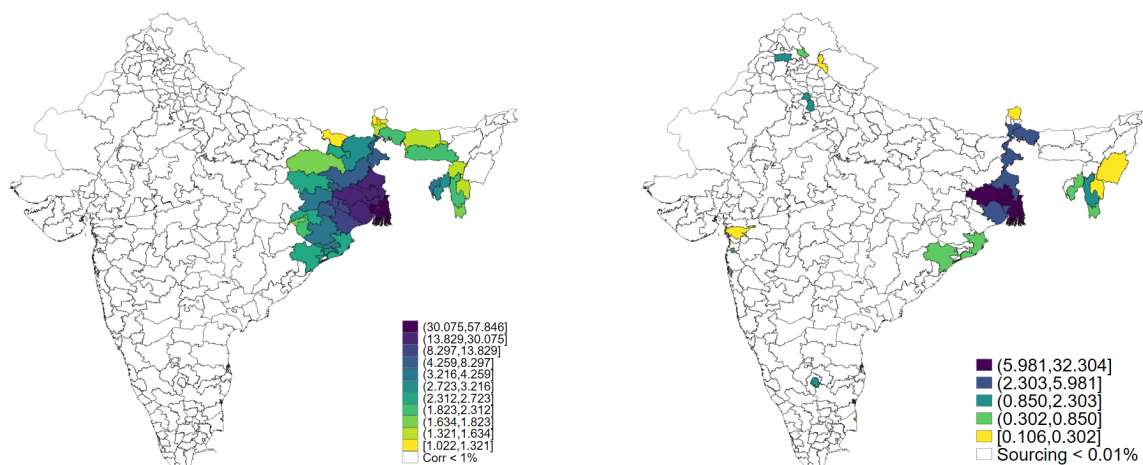
Figure 8: Shock probabilities and wages



Note. In this figure, we plot model-derived nominal (left panel) and real wages (right panel) against the estimated shock probabilities. Figure D3 further plots the price index and the variance in real wages against the shock probabilities.

closer areas than, say, the far south of India. Firms also source from districts less spatially correlated with their own. Notice that the sourcing patterns include several zeros in equilibrium.

Figure 9: Spatial correlation and sourcing shares: Kolkata

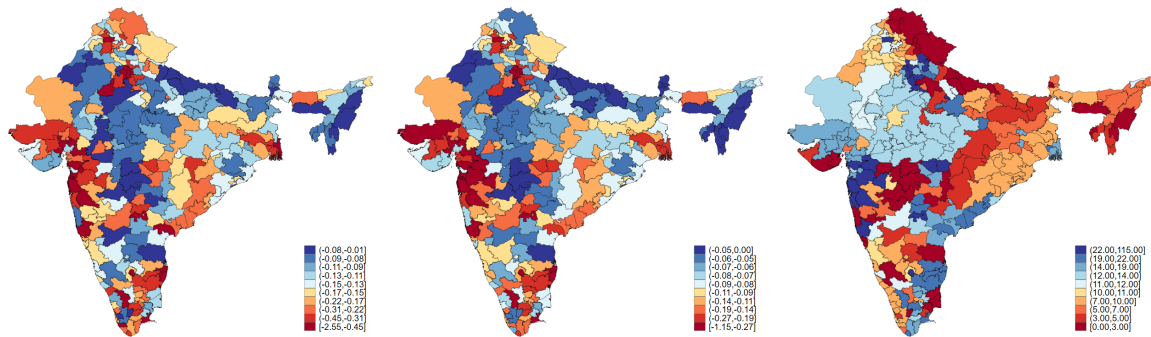


Note. In this figure we plot the estimated spatial correlation in disruption probabilities with other districts for Kolkata (left panel) and the sourcing shares of Kolkata district with all other districts (right panel).

Shock propagation Our framework can also be used to assess the effect of disruptions ex-post for aggregate welfare. In Panel A of Figure 10, we show, for each origin district, the impact of a disruption in that district on the real wages of all other dis-

tricts (including itself). We use the size of the labor force in each district to compute the weighted average of the effect. The impact of a realized disruption in a district on the rest depends on the affected district’s importance as a supplier. The effects vary widely by district, with shocks that materialize in lower-risk or more productive districts that are more important as sourcing locations having larger welfare consequences. While the “own” effect of the shock is important, a large component (62.6% on average) of the aggregate welfare changes happens through the propagation of the shock (Panel B of the figure). Here, we plot the aggregate welfare changes caused by the incidence of a disruption in each origin district, excluding the own effect. Finally, Panel C illustrates the number of districts that experience a welfare decline when an origin district experiences a disruption.

Figure 10: Shock Propagation



(a) Weighted Average Welfare Change (b) Weighted Average Welfare Change in Other Regions (c) Number of regions with Welfare Decline > 1%

Note. In Panel A, for each district, we compute the impact a materialized disruption has on the real wages of all other districts (including itself). We then use the labor force in each district to calculate the weighted average of the impact. Panel B removes the own-impact in real wages of a disruption to isolate the “propagation” effect to other districts. Panel C reports the number of districts that experience a welfare decline when the district experiences a disruption.

4.4 Trade Counterfactuals

We compare welfare in the calibrated model to regional autarky and free trade. Throughout, we decompose the welfare effects on the changes in expected real wages and their volatility, capturing the first- and second-moment effects in the model.

Expected welfare under baseline and autarky. The comparative statics in Section 3 show that with identical regional fundamentals, calibrated trade costs, and

independent disruption probabilities, expected real wages are lower for all regions with costly trade than in autarky, and their variance is also lower. To assess the quantitative relevance of this mechanism in the calibrated model with varying regional fundamentals, estimated trade costs, and disruption probabilities that are spatially correlated, we compute the difference in expected real wages in the baseline model with the model-implied expected real wages given the same regional fundamentals, disruption probabilities, and infinite trade costs.

Figure 11 illustrates the spatial variation of expected real wages in the baseline model and the autarky counterfactual. On average, expected real wages are 3.1% higher in autarky than in the baseline model. The variance of real wages is 9.2% higher in autarky, validating the quantitative relevance of the main comparative statics exercises. Overall, autarky is welfare decreasing for all regions. Welfare decreases on average by 7.3%, as the change in volatility more than offsets the gain in log expected real wages. 0.74% of districts see real wage declines, unlike in the comparative statics where all regions had higher real wages in autarky.

Expected welfare under baseline and free trade. In contrast, Figure 11 shows that expected real wages are higher for all regions under a free trade counterfactual, *and* their volatility is lower, so the welfare gains from free trade are large. To implement free trade in our quantitative exercise, we set the iceberg trade costs to 1 between all districts. Under free trade, expected real wages are, on average, 5.9% higher than in the baseline, whereas the variance of real wages is 2.8% lower. Welfare is on average 8.9% higher and no district is worse off under free trade.

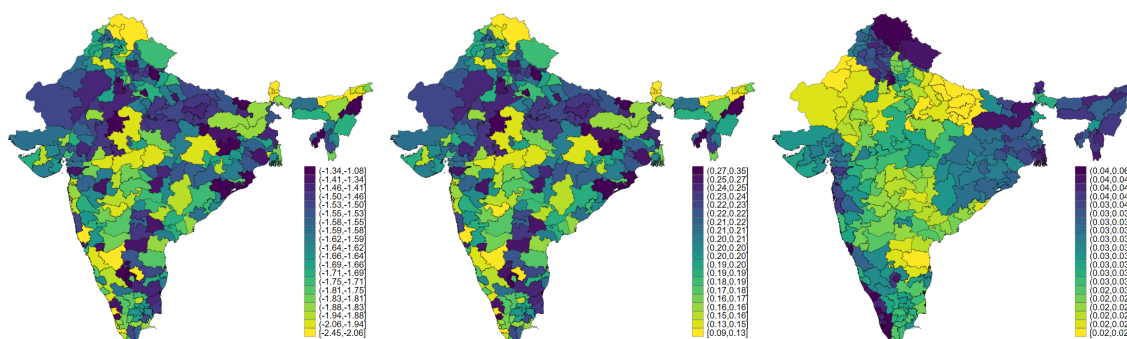
4.5 Climate Change Counterfactuals

We next study the implications of varying climate risk in our model. We estimate the share of our model-implied shock probabilities that can be explained by climate-risk-related variables such as rainfall or flooding events. Through the lens of our model, these probabilities capture the risk firms assign to each district. However, as discussed above, the risk associated with each region can be due to climate risk, as well as other regional characteristics such as infrastructure or governance. In this section, to highlight the implications of changing climate risk, we hold all other sources of risk constant and change only the climate risk of each region relative to the baseline.

To discipline how climate risk changes, we proceed as follows: First, we regress

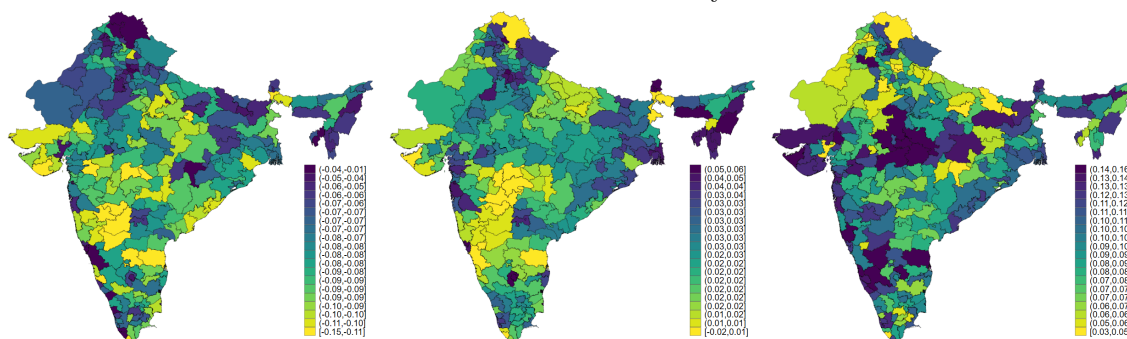
Figure 11: Quantitative results

Panel A: Baseline



(a) Expected Welfare (b) Log Expected Real Wages (c) Variance of Real Wages

Panel B: Δ in Autarky



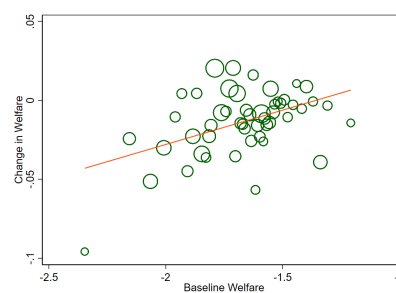
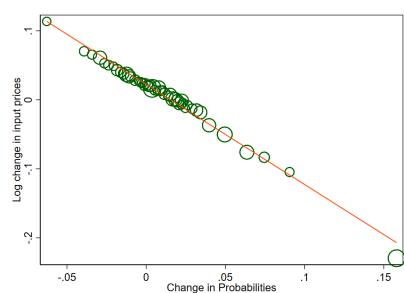
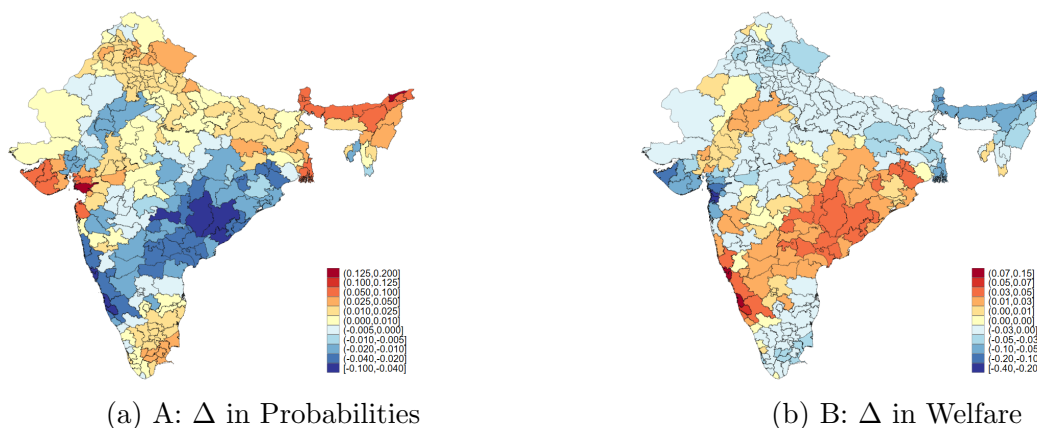
(d) Expected Welfare (e) Log Expected Real Wages (f) Variance of Real Wages

Note. Panel A shows welfare, expected real wages and their variance in the baseline calibrated model. Panel B shows percentage changes in these variables under the autarky counterfactual relative to the baseline scenario. The maps for change under free trade can be found in Appendix D5

the inverse logit transformation of our probabilities on historical measures of rainfall, coastal flooding, riverine flooding, temperature, and the SPI presented in Figure 6. Second, we use the estimated coefficients, shown in Column 1 of Table 4, to predict the counterfactual disruption probabilities in 2050 for our five climate measures, while holding constant the unexplained variation in these probabilities. This method yields how the probabilities would change if climate variables evolve as predicted in the RCP 4.5 scenario of the International Project for Climate Change (IPCC).

Panel A of Figure 12 illustrates how these probabilities change across space in our main counterfactual. As the figure makes clear, there is wide variation in the changes in climate risk, with the northeast and parts of the west coast seeing large increases

Figure 12: Counterfactuals: Climate Risk Increase



Note. In this figure, we plot the change in probabilities of climate risk (panel A), the change in welfare (panel B), the relationship between the change in input prices and changes in probabilities change in expected real wages (panel C), and the relationship between the change in welfare in the counterfactual and the welfare at baseline when climate risk increases as described in Section 4.5.

in risk, while the central part of the country sees decreases in risk. On average, risk increases by 1.1 percentage points. Panel B illustrates the change in expected welfare in this counterfactual. Welfare on average decreases by 2.01%. There is wide spatial variation, with a range of 3.11pp, and some of the less risky regions see welfare gains. 62.73% of districts see real wage declines.

To understand the mechanisms at work, Panel C shows how the change in district supplier prices correlates with changes in district risk. Input prices offered by intermediate firms from the district decrease the most for districts experiencing the largest increases in risk. This negative terms-of-trade effect arises from the decline in nominal wages in these risky regions in equilibrium.³³

³³Recall input prices $p_i = \frac{w_i}{z_i}$. Effectively, the nominal wages in risky regions are decreasing, by *more* than the increase in risk as firms diversify away from riskier regions.

Panel D illustrates the change in welfare, and relates it to the initial district welfare. This highlights the distributional consequences of climate change in our quantification: the change in welfare is positively correlated with initial welfare. In other words, initially well-off regions see relative welfare improvements following climate risk increases, while initially worse-off regions see welfare declines. A key quantitative finding is that here, sourcing diversification of firms amplifies the effects of climate risk increases. Climate risk will not only subject riskier regions to increased shocks, but also to a decrease in real wages as firm supply chains become less reliant on these regions.³⁴ Table 5 summarizes the quantitative results across counterfactuals.

Table 5: Model Counterfactuals: Summary

Counterfactual	Δ in Welfare		Δ in log Expected Real Wages		Δ in Real Wage Volatility		% districts Real wage declines
	Avg. change	Range	Avg. change	Range	Avg. change	Range	
<u>Baseline risk</u>							
Autarky	-7.29	2.92	3.10	1.87	9.25	3.99	0.74%
Free Trade	8.94	2.30	5.92	1.70	-2.84	0.96	0.00%
<u>Alternative risk</u>							
Climate change	-2.01	3.11	-1.96	3.10	0.15	0.13	62.73%
Δ in Rainfall and Flood Risk Only	-0.24	3.52	-0.25	3.46	0.06	0.13	25.09%
Δ in Temperature and SPI Risk Only	-1.76	2.69	-1.72	2.64	0.06	0.13	86.72%

Note. This table shows statistics of the distribution of percentage changes between the baseline scenario with current climate risk and costly trade and other scenarios, weighted by district population. Range refers to the interquartile range.

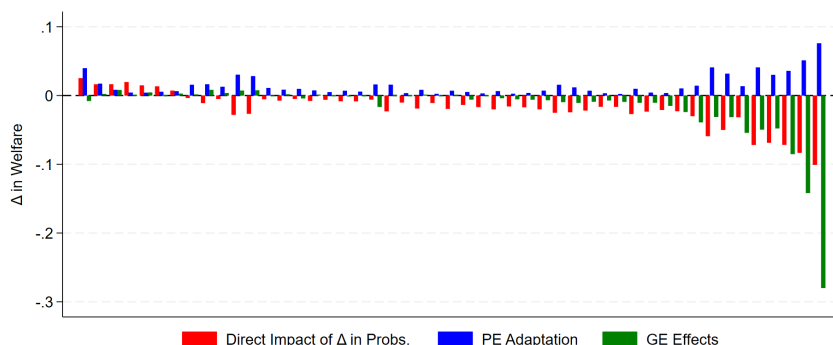
Decomposing the effects of climate change adaptation. As a final exercise, we decompose the change from our baseline economy to the counterfactual economy with increased climate risk into three components

$$\begin{aligned}
 \Delta \mathcal{W}_i = & \underbrace{\mathcal{W}_i(G', \mathbf{M}_i(G', \mathbf{w}')) - \mathcal{W}_i(G', \mathbf{M}_i(G', \mathbf{w}))}_{\text{G.E. Effect}} \\
 & + \underbrace{\mathcal{W}_i(G', \mathbf{M}_i(G', \mathbf{w})) - \mathcal{W}_i(G', \mathbf{M}_i(G, \mathbf{w}))}_{\text{P.E. adaptation}} + \underbrace{\mathcal{W}_i(G', \mathbf{M}_i(G, \mathbf{w})) - \mathcal{W}_i(G, \mathbf{M}_i(G, \mathbf{w}))}_{\text{Direct effect of climate change}},
 \end{aligned} \tag{26}$$

where X' refers to the changed climate risk scenario. The direct effect captures the effect of changing climate risk, without firm adaptation. In practice, we start in the baseline equilibrium, but simulate a model where shocks are drawn from the new distribution with changed climate risk. Agents' beliefs in this step of the decomposition are therefore not rational, and we refer to this as climate “myopia”. The P.E.

³⁴This is not a mechanical result, but rather, depends on the spatial distribution of climate risk, the initial equilibrium, and the IPCC predictions for which areas get riskier. If initially higher welfare areas saw larger changes in predicted climate risk, they would not see relative welfare increases.

Figure 13: Distributional Implications of Climate Change Adaptation



Note: This figure plots the terms in 26, binning regions into 50 bins. The x-axis orders regions by their change in disruption probabilities. The red bars show direct effects, blue bars show the P.E. effects and green bars show the G.E. effects.

effect considers the effect of firm adaptation to the climate risk in partial equilibrium, holding all prices fixed. Finally, the general equilibrium effects allow for equilibrium price adjustment.

Figure 13 contains the results. The red bars show the direct effects of changing climate risk, which are heterogeneous across regions. The blue bars have the welfare effects of partial equilibrium adaptation to new risk. Holding prices fixed, such adaptation is always beneficial, even for regions with increased risk. For some districts the P.E. term offsets the increased direct risk. The green bars show the general equilibrium effects on prices. Regions facing the largest increases in disruption risk experience significant welfare declines due to general equilibrium price adjustments. As firms across all regions reduce demand for their inputs, wages fall, compounding welfare losses beyond the direct impact of rising risk. These regions fare worse than they would if firms were myopic and did not to adapt to the heightened risks.

Robustness and extensions. In addition to our main climate counterfactual, we also consider scenarios where only rainfall and flood risk, or only temperature changes and SPI changes occur. Table 5 summarizes the results. While in both cases average welfare declines and there is wide spatial variation, under the scenario of only temperature/SPI changes, 86.72% of districts see real wage declines, while with only rainfall/flood risk increases, 25.09% of districts see real wage declines.

Appendix E estimates two alternative models and conducts the counterfactuals in

these cases. We first consider a model where the district probabilities are obtained from projections on climate-related variables, as discussed in Section 4.2. Second, we consider a model where the input bundle is CES and so features love-for-variety effects, with a substitution elasticity of 3.1 (Peter and Ruane, 2023). Table E3 summarizes the results for these two alternative models. Strikingly, the two models deliver very similar implications for the climate counterfactuals, in terms of the welfare declines and spatial variation.

Appendix C.4 shows that mean and median product-level inventories are less than a month’s sales and that inventories are not correlated with the prevalence of multi-sourcing. While inventories and multisourcing would appear to be alternative strategies for risk mitigation, in our data it appears firms are systematically choosing multi-sourcing. We note that our calibrated model *without* inventories implies that sales declines less than inputs upon the incidence of a shock. Equation (10) illustrates that the partial elasticity of firm profits to delivered inputs is $\frac{(1-\beta)(\sigma-1)}{\beta+\sigma(1-\beta)}$. Quantitatively, given our parameter calibration, this implies that sales fall by 47% of the fall in inputs, which is very similar to the 44% drop observed in the event-studies.

5 Conclusion

Climate risk is an increasingly important concern worldwide, with large projected economic impacts. Adaptation of firm supply chains to perceived climate risk is a crucial channel through which economies might adjust to climate risk. This paper provides empirical evidence suggesting firm supply chains are structured taking climate risk into account. Our new model of firm supply chain decisions under risk incorporates key patterns we see in the data – firms face a trade-off between lower risk and higher input costs. Quantitative results from the model show that, on the one hand, input sourcing decisions mitigate the effects of climate risk on welfare, especially through a reduction of the volatility of output, as firms diversify their suppliers. Yet, on the other hand, they amplify the distributional effects, as regions that face increasing climate risks will also suffer lower real wages from the general equilibrium effects of firm adaptation.

References

- Akerberg, D. A., K. Caves, and G. Frazer (2015, November). Identification Properties of Recent Production Function Estimators. *Econometrica* 83(6), 2411–2451.
- Adamopoulos, T. and F. Leibovici (2024). Trade Risk and Food Security. *Mimeo*.
- Alfaro Ureña, A., M. Fuentes Fuentes, I. Manelici, and J. P. Vasquez (2018). Costa Rican Production Network: Stylized Facts. *Working paper*.
- Allen, T. and D. Atkin (2022). Volatility and the Gains from Trade. *Econometrica* 90(5), 2053–2092.
- Antràs, P., T. C. Fort, and F. Tintelnot (2017, September). The Margins of Global Sourcing: Theory and Evidence from US Firms. *American Economic Review* 107(9), 2514–64.
- Balboni, C. (2025). In Harm’s Way? Infrastructure Investments and the Persistence of Coastal Cities. *American Economic Review* 115(1), 77–116.
- Balboni, C., J. Boehm, and M. Waseem (2024). Firm Adaptation and Production Networks: Structural Evidence from Extreme Weather Events in Pakistan. *Mimeo*.
- Barrot, J.-N. and J. Sauvagnat (2016). Input Specificity and the Propagation of Idiosyncratic Shocks in Production Networks. *Quarterly Journal of Economics* 131(3), 1543–1592.
- Bilal, A. and E. Rossi-Hansberg (2023, June). Anticipating Climate Change Across the United States. Working Paper 31323, National Bureau of Economic Research.
- Blaum, J., F. Esposito, and S. Heise (2024). Input Sourcing Under Risk: Evidence from U.S. Manufacturing Firms. *Mimeo*.
- Boehm, C. E., A. Flaaen, and N. Pandalai-Nayar (2019, March). Input Linkages and the Transmission of Shocks: Firm-Level Evidence from the 2011 Tōhoku Earthquake. *The Review of Economics and Statistics* 101(1), 60–75.
- Boehm, C. E., A. A. Levchenko, and N. Pandalai-Nayar (2023, April). The Long and Short (Run) of Trade Elasticities. *American Economic Review* 113(4), 861–905.

- Caliendo, L. and F. Parro (2015, 11). Estimates of the Trade and Welfare Effects of NAFTA. *The Review of Economic Studies* 82(1), 1–44.
- Carvalho, V. M., M. Nirei, Y. U. Saito, and A. Tahbaz-Salehi (2021, May). Supply Chain Disruptions: Evidence from the Great East Japan Earthquake. *Quarterly Journal of Economics* 136(2), 1255–1321.
- Caselli, F., M. Koren, M. Lisicky, and S. Tenreyro (2019, 09). Diversification Through Trade. *The Quarterly Journal of Economics* 135(1), 449–502.
- Castro-Vincenzi, J. (2024, June). Climate Hazards and Resilience in the Global Car Industry. *mimeo*.
- Cruz, J.-L. and E. Rossi-Hansberg (2024). The Economic Geography of Global Warming. *Review of Economic Studies* 91(2), 899–939.
- Desmet, K., R. E. Kopp, S. A. Kulp, D. K. Nagy, M. Oppenheimer, E. Rossi-Hansberg, and B. H. Strauss (2021, April). Evaluating the Economic Cost of Coastal Flooding. *American Economic Journal: Macroeconomics* 13(2), 444–86.
- Esposito, F. (2022). Demand risk and diversification through international trade. *Journal of International Economics* 135(103562).
- Farrokhi, F. and A. Lashkaripour (2024). Can Trade Policy Mitigate Climate Change? . Working Paper.
- Grossman, G., E. Helpman, and H. Lhuillier (2023). Supply Chain Resilience: Should Policy Promote Diversification or Reshoring? *Journal of Political Economy* 131.
- Grossman, G., E. Helpman, and A. Sabal (2024). Resilience in vertical supply chains. *Mimeo*.
- Gu, G. and G. Hale (2023, December). Climate Risks and FDI. *Journal of International Economics* 146.
- Helpman, E. and A. Razin (1978, August). Welfare Aspects of International Trade in Goods and Securities. *The Quarterly Journal of Economics* 92(3), 489–508.
- Hsiao, A. (2023). Sea Level Rise and Urban Adaptation in Jakarta. *Mimeo*.

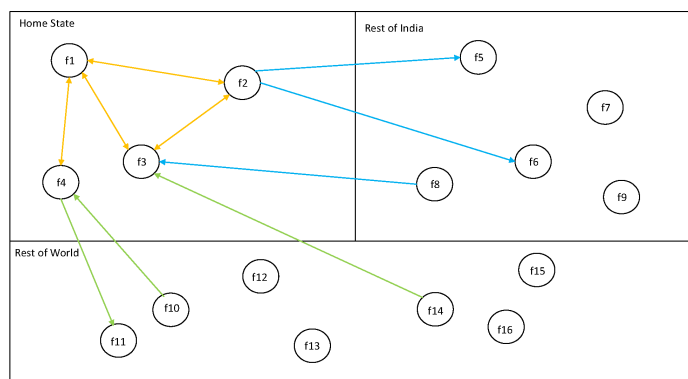
- Huo, Z., A. A. Levchenko, and N. Pandalai-Nayar (2024, 03). International Comovement in the Global Production Network. *The Review of Economic Studies*.
- Indaco, A., F. Ortega, Taspınar, and Süleyman (2020, 11). Hurricanes, Flood Risk and the Economic Adaptation of Businesses. *Journal of Economic Geography* 21(4), 557–591.
- Jia, R., X. Ma, and V. W. Xie (2022, July). Expecting Floods: Firm Entry, Employment, and Aggregate Implications. Working Paper 30250, National Bureau of Economic Research.
- Johnson, R. C. and G. Noguera (2012). Accounting for Intermediates: Production Sharing and Trade in Value Added. *Journal of International Economics* 86(2), 224 – 236.
- Khanna, G., N. Morales, and N. Pandalai-Nayar (2025). Supply Chain Resilience: Evidence from Indian Firms. NBER Working Papers 30689, National Bureau of Economic Research, Inc.
- Kopytov, A., B. Mishra, K. Nimark, and M. Taschereau-Dumouchel (2024, September). Endogenous Production Networks under Supply Chain Uncertainty. *Econometrica*.
- Korovkin, V., A. Makarin, and Y. Miyauchi (2024). Supply Chain Disruption and Reorganization: Theory and Evidence From Ukraine’s War. Working Paper.
- Nath, I. (2024). Climate Change, the Food Problem, and the Challenge of Adaptation through Sectoral Reallocation. *Mimeo*.
- Pankratz, N. and C. Schiller (2024). Climate Change and Adaptation in Global Supply-Chain Networks. *Review of Financial Studies* 37(6), 1729–1777.
- Pellet, T. and A. Tahbaz-Salehi (2023). Rigid Production Networks. *Journal of Monetary Economics* 137(C), 86–102.
- Peter, A. and C. Ruane (2023, May). The aggregate importance of intermediate input substitutability. *NBER Working Paper No. 31233*.
- Yi, K.-M. (2003, February). Can Vertical Specialization Explain the Growth of World Trade? *Journal of Political Economy* 111(1), 52–102.

Supplemental Online Appendix

A Details on the Firm-to-Firm Data

We illustrate a stylized example of our establishment-level networks data in Figure A1. As the diagram shows, we observe all transactions where one node of the transaction is within the state. This includes all transactions between establishments within the state (the yellow lines), any inflows from or outflows to the rest of the country (the blue lines), and all international imports and exports (the green lines).

Figure A1: Stylized Example of Establishment-Level Network



Notes: Stylized example of establishment-level data. The upper half represents the country, and upper left quadrant represents the state in question. The data includes all transactions within the state, and all transactions where one node of the transaction (either buyer or seller) is in the state.

The data report value and quantity of traded items, so we can construct unit values. To do this, we aggregate values and quantities at the four-digit HSN/month/transaction level, and then construct implied unit values. We can then collapse the data at the 4-digit HSN/month level to construct average unit values, the number of transactions between each seller and buyer pair, and the total value of the goods transacted. This is the foundation of the firm-to-firm dataset we use in the analysis. Additionally, we can aggregate these data to the buyer level, which we use in our reduced-form section. Table A1 summarizes our primary variables of interest using this dataset. In Table A2 we present statistics on the number of buyers per supplier and suppliers per buyer. Despite differences in region sizes, the distribution of firms follows closely the one documented by [Alfaro Ureña et al. \(2018\)](#) for Costa Rica.

Table A1: Summary Statistics for Main Variables

Outcome	Mean	p25	p50	p75
Separation Rate (%)	30.9	0	16.67	52.78
Entry Rate (%)	74.06	0	50	106.67
Net Separations (%)	-43.12	-70	0	0
Real Input Value (log)	14.91	12.48	14.55	16.96
Real Sales (log)	16.33	13.57	16.05	18.66
Avg. Supplier Size (millions of rupees)	106.42	9.65	34.04	127.49
Avg. Supplier Outdegree	43.04	3.3	10.97	31.99
Share Purch. Lgst. Supplier (%)	52.39	31.06	47.84	71.82
Number Products	12.05	3	7	14
Share Purch. Diff. Prod. (%)	60.19	21.25	72.78	97.81
Supply Chain Depth	32.32	28.15	31.46	36.35
Number Suppliers	12.35	3	7	14
Avg. Distance (km)	486.71	97.13	251.65	712.75
Share Purch. Non-Home State (%)	38.54	0	24.42	78.48

Note. Summary statistics for key outcomes to describe the network calculated in December 2019-February 2020. Number of firms included in calculations: 136,562.

Table A2: Distribution of buyers and suppliers

	Mean	SD	10th	25th	50th	75th	90th	95th	99th
N suppliers per buyer	8.0	23.6	1	1	3	8	18	29	72
N of buyers per supplier	16.3	55.3	1	1	4	12	36	65	194
N supplier districts per buyer	3.5	4.4	1	1	2	4	7	11	21
N buyer districts per supplier	3.1	3.0	1	1	2	4	7	10	14

Note. We calculate network characteristics for the year 2019. The top two rows compute the number of buyers per supplier and suppliers per buyer. The bottom rows compute the number of supplier districts per buyer and number of buyer districts per supplier.

B Empirical Facts Appendix

In Table 1 we show that firms seem to multi-source products even within detailed product categories. We proceed to show that such results is not driven by retailers and wholesalers. While we cannot directly identify retailers and wholesalers in our data, we can use the pattern of their transactions to infer firms that likely belong to those industries. For retailers, we expect that they would sell their goods predominantly to final consumers instead of shipping their goods to other firms. Hence, they should show up as having zero sales in our data. For wholesalers, we expect that they would not transform the products they buy in order to sell them. Hence, we identify

as wholesaler firms that buy and sell the same HSN-4 products. Of course, these classifications will be overestimating retailers and wholesalers, as manufacturing firms might buy and sell the same 4-digit product or not ship goods to other firms. However, we want to corroborate that our results are robust to excluding these firms.

From our sample in 2019, we have a total of 195,872 firms. Of those, 7,867 fall under our classification of wholesalers and 137,574 fall under our classification of retailer. As shown in Table B1, the distributions of regions sourced from stay fairly constant when excluding such firms. Similarly, we show that the results are consistent when we look at the number of suppliers per product as opposed to supplier districts. As shown in Table B2, there is a slightly larger fraction of firms that source from more than one supplier than when looking at sourcing from different supplier districts.

Table B1: Firms that source from multiple districts (excluding wholesale and retail)

Number of districts	Share of buyers		Share of buyers x HSN-4		Share of buyers x HSN-8	
	Firms	Value	Firms	Value	Firms	Value
1	12.0%	1.5%	69.4%	9.6%	80.9%	21.6%
2	13.4%	2.0%	16.4%	9.8%	12.6%	15.4%
3	12.4%	2.8%	6.4%	8.9%	3.5%	11.1%
4	10.6%	3.0%	3.0%	8.0%	1.4%	10.0%
5	8.9%	3.1%	1.6%	7.5%	0.6%	5.9%
6	7.3%	3.7%	1.0%	6.4%	0.3%	4.5%
7	6.0%	3.4%	0.6%	4.5%	0.2%	4.5%
8	5.0%	3.8%	0.4%	4.6%	0.1%	2.9%
9	4.1%	3.4%	0.3%	4.0%	0.1%	3.1%
10+	20.3%	73.3%	0.9%	36.6%	0.2%	21.1%

Note. Column 1 aggregates the data at the firm level and computes the share of firms that source from a certain number of districts. Column 2 calculates the fraction of total value purchased by number of supplier districts sourced from. Columns 3-4 aggregate the data at the firm-by-4-digit product level, and Columns 5-6 at the firm-by-8-digit product level. We exclude likely-retailers and likely-wholesalers from the analysis.

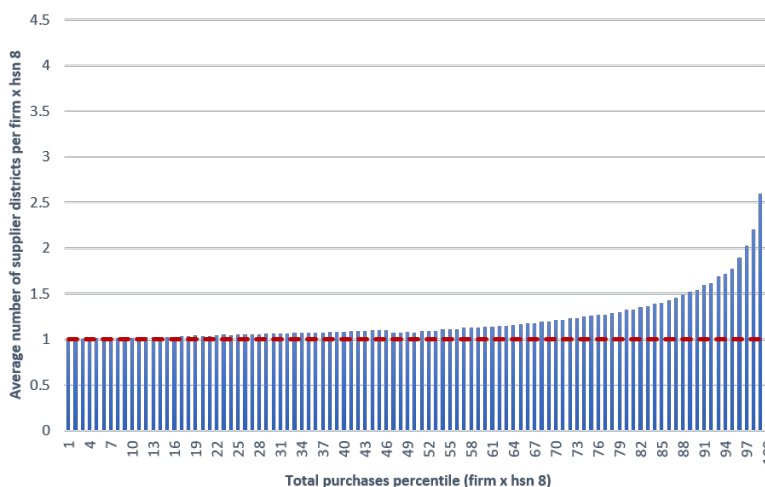
Next, we show that firms that have larger purchases of a given product are more likely to source from multiple regions. To see this, we rank all firm-by-8-digit HSN pairs into percentiles based on total purchases, where the higher percentiles include the firm-product pairs with the higher purchase volume. As shown in Figure B1, the smallest firm-product pairs tend to only source from a single supplier. However, towards the end of the distribution, the largest firm-product pairs source, on average, from more than one region. Firms above the 95th percentile source, on average, from two districts, and firms in the top percentile source from four. This suggests that larger, more productive firms are more likely to multisource.

Table B2: Share of firms that source from multiple suppliers

Number of suppliers	Share of buyers		Share of buyers x HSN-4		Share of buyers x HSN-8	
	Firms	Value	Firms	Value	Firms	Value
1	28.6%	2.3%	70.6%	15.7%	81.7%	24.5%
2	16.0%	6.7%	15.0%	9.2%	11.7%	17.1%
3	10.4%	2.3%	5.7%	7.2%	3.4%	12.7%
4	7.4%	2.2%	2.9%	5.8%	1.4%	8.4%
5	5.4%	1.8%	1.6%	6.5%	0.7%	5.2%
6	4.2%	1.8%	1.0%	4.8%	0.4%	4.7%
7	3.4%	1.8%	0.7%	3.5%	0.2%	3.0%
8	2.7%	1.6%	0.5%	2.9%	0.1%	2.6%
9	2.3%	2.0%	0.4%	3.3%	0.1%	2.5%
10+	19.6%	77.4%	1.6%	41.1%	0.3%	19.4%

Note. In this table we look at number of supplier firms instead of number of supplier districts. Column 1 aggregates the data at the firm level and computes the share of firms that source from a certain number of suppliers. Column 2 calculates the fraction of total value purchased by number of suppliers sourced from. Columns 3-4 aggregate the data at the firm-by-4-digit product level, and Columns 5-6 at the firm-by-8-digit product level.

Figure B1: Number of supplier districts by total purchases



Note. We rank all firm-product pairs into percentiles (1-100) based on the volume of total purchases in 2019. For each percentile (in the horizontal axis), we compute the average number of districts the firm-product pairs source from.

However, firm size does not drive the descriptive patterns shown in Figures 1a-1c. In Table B3, we document that our descriptive patterns under Fact 2 are not driven by firm size, product composition or capacity of suppliers. We run a regression at the product-firm level as shown in equation 27.

$$\log y_{j,p} = \beta \mathbb{1}(\text{Firm } j \text{ multisources } p) + \gamma X_j + \delta_p + \epsilon_{j,p} \quad (27)$$

where $\log y_{j,p}$ is the log of the average characteristic of a firm’s suppliers such as average distance to suppliers, rainfall of supplier districts, riverine flooding of supplier districts and prices paid to suppliers. The key explanatory variable here is $\mathbb{1}(\text{Firm } j \text{ multisources } p)$ which is a dummy that indicates whether the firm sources product p from more than one district. Importantly, we control for product fixed effects, the log of total purchases by firm j , and the log average sales of suppliers.

Table B3 shows that our descriptive patterns are robust to adding these controls. Multi-sourcers buy products from distances 76% farther than single-sourcers. They source from districts with 2.3% lower rainfall and 1.4% lower river flooding levels. Finally, they pay 44% higher input prices than single sourcers. Product fixed effects help rule out that the differences between single and multi sourcers are driven by differences in product quality. The own purchases control rules out that the patterns are driven by differences in firm size (e.g. large firms multisource more and pay higher prices). Finally, the control for supplier size helps us rule out that the reason for multisourcing is that suppliers don’t have enough capacity to meet demand.

Table B3: Supplier characteristics by number of districts sourced from

	Log (Distance to suppliers)	Log(Daily Rainfall)	Log(Historical riverine flooding)	Log(Price of inputs)
$\mathbb{1}(\text{Multisourcer})$	0.760***	-0.0229***	-0.0140***	0.441***
N	739,520	739,520	739,520	739,520
R-sq	0.327	0.271	0.124	0.545

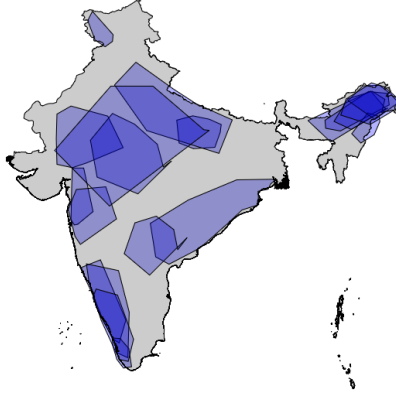
Note. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. We run a cross-sectional regression at the firm (j), 8-digit product (p) level. The outcome is the log average distance to suppliers (column 1), log average daily rainfall at suppliers’ district (column 2), log average riverine flooding at suppliers’ district (column 3) and log average price of inputs (column 4). The main regressor is a dummy variable on whether the firm sources the HSN-8 product from more than one district. All regressions include HSN-8 product fixed effects and controls for log size of the firm and log average size of suppliers.

B.1 Responses to flooding events - additional results.

B.1.1 Climate Data

We use data from the Dartmouth Flood Observatory to identify geocoded flooding events throughout India for our event study analysis. As shown in Figure B2, we identify 19 events of large monsoonal floods throughout India between 2018 and

Figure B2: Monsoonal rain floods, 2018-2020



Note: The figure plots the geographic coverage of all large floods that occurred between 2018 and 2020, as described in the Dartmouth Flood Observatory.

2021. For our event study analysis, we limit the set of floods to those that occurred outside of our state, for which we have at least 3 months of data before and after the flood, and where at least 200 buyers in our state transacted with affected suppliers the period before the flood. These restrictions leave us with seven large flood events, which we use in our analysis Fact 4 in Section 2.2.

B.1.2 Two-way Fixed Effects Analysis

In this Section, we delve deeper into the event-study results presented in Figure 2. For expositional clarity, we run difference-in-difference specifications which we summarize in Table B4. In the top panel, we present the two-way fixed effects specifications with continuous treatment as described below. In the middle panel, we present the results for a binary treatment. Finally, in the bottom panel, we present the results using the Local Projections Difference-in-Differences (LP-DID) estimator developed by Dube et al. (2023). This last set of estimates further accounts for issues raised by recent discussions on two-way fixed effects methods. We begin with documenting the direct effect on suppliers in flood-hit zones, where we examine outcomes $y_{j,t,k,\tau}$ for firm j , in period t , industry k , and event τ .

$$y_{j,t,k,\tau} = \alpha \mathbb{1}(\text{Exposed to flood})_{j\tau} \times \text{Post}_{t,\tau} + \sum_{x=-5}^{x=+5} [\delta_{\tau,x} + \beta_x X_{j,x<0}] + \delta_j + \delta_{k,t} + \epsilon_{j,t,k,\tau}. \quad (28)$$

Here, “Exposed to flood $_{j\tau}$ ” takes a value of 1 if firm j was exposed to a particular flood. We index the months before and after flood happened by x , with $x = 0$ being

the month the flood τ occurs. We include a wide range of high-dimensional fixed effects to account for confounding shocks. These include firm fixed effects δ_j that control for firm-specific time-invariant differences; industry-by-time fixed effects $\delta_{k,t}$ that control for industry-specific shocks; and flood event-time since flood fixed effects $\delta_{\tau,x}$ that control for aggregate trends around the flood event that affect all firms (including those not in the flood-exposed areas). We also control for firm size-specific shocks, by controlling for sales in the five months before the flood $X_{j,x<0}$, interacted with time-since flood indicators. In all difference-in-difference results we restrict the post period to 3 months after the flood. Consistent with the results in Figure 2a, we find that affected sellers experience a 13% decline in sales, on average, with respect to non-affected firms the three months after the flood occurs.

Next, we look into the effect of a flood for buyers located in our state. We use the existing supplier network (in the 5-months leading to the flood) as a measure of the exposure to the disruption, as described in equation B.1.2.

$$(\text{Supplier Exposure})_{j\tau} = \sum_i^N s_{i,j,\tau,x<0} \times \mathbb{1}(\text{Supplier } i \text{ exposed to flood in } \tau) ,$$

where $s_{i,j,\tau,x<0}$ is the value of purchases that firm j buys from firm i , relative to firm j 's total purchases, over the five months before the flood. We then standardize this index and interact it with a post flood indicator to study how buyers were affected when their suppliers were hit. We examine outcomes $y_{j,t,k,\tau}$ for firm j , in period t , and industry k , measured in event-time (since flood) τ equation 29:

$$y_{j,t,k,\tau} = \gamma (\text{Flood Exposure})_{j\tau} \times \text{Post}_{t,\tau} + \sum_{x=-5}^{x=+5} [\delta_{\tau,x} + \beta_x X_{j,x<0}] + \delta_j + \delta_{r,k,t} + \epsilon_{j,t,k,\tau}. \quad (29)$$

The fixed effects are similar to equation 28 but we add an industry-region-time fixed effect $\delta_{k,r,t}$ to control for local demand shocks affecting the region-industry of the firm. In columns 2-4 of Table B4, we present the results for the outcomes of log total purchases (column 2), log purchases of returning suppliers (column 3) and log purchases of new suppliers (column 4). Returning suppliers are those who transacted with the firm within 3-months before the shock, and we track the purchases from that set of suppliers throughout time. New suppliers are defined as suppliers who transact with the firm in a given period who have not transacted before. Difference-in-difference results are consistent with the event studies in Figures 2b and 2c.

In column 5 of Table B4, we present the results for the outcome of buyer sales.³⁵ Buyers with one standard deviation higher exposure experience a decline in sales of 2% relative to buyers with average exposure. When considering the binary treatment in the second panel, we find that firms exposed to the flood through their suppliers experience a decline of 7% relative to firms that are not exposed.

When focusing on the result for sales with binary treatment, we find that for every 1% decrease in the purchases for exposed buyers, sales drop by 0.44% (exposed buyers decrease purchases by 16% and sales by 7% relative to non exposed buyers). We compare this result to the one implied by our model which, using a back-of-the-envelope calculation, indicates that for every 1% decline in purchases, sales decrease by 0.47%. The close result is reassuring given that our sales result is untargeted by our estimation process.

Table B4: Regression results on the impact of floods.

	Supplier sales	Buyer Purchases - Total	Buyer Purchases - Returning Suppliers	Buyer Purchases - New suppliers	Buyer Sales	Input prices
Continuous treatment						
Standardized exposure $\times \mathbb{1}(\tau \geq 0)$	-	-0.05*** (0.003)	-0.05*** (0.003)	-0.03*** (0.01)	-0.02*** (0.01)	-0.009 (0.01)
<i>N</i>	-	1,218,663	1,160,881	606,655	468,280	1,912,563
Binary treatment						
$\mathbb{1}(\text{Exposure} \geq 0.1) \times \mathbb{1}(\tau \geq 0)$	-0.13*** (0.02)	-0.16*** (0.01)	-0.24*** (0.01)	-0.13*** (0.02)	-0.07** (0.03)	0.004 (0.01)
<i>N</i>	1,604,955	1,218,663	1,160,881	606,655	468,280	1,912,563
Local projections with binary treatment						
Difference between pooled pre and post period	-0.29*** (0.06)	-0.17*** (0.02)	-0.12*** (0.03)	-0.15** (0.08)	-0.03 (0.06)	0.007 (0.09)
<i>N</i>	742,966	897,777	829,534	130,600	413,392	716,388

Note. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Column 1 presents the estimates for equation 28 for suppliers affected by the floods. Columns 2-5 present the results for equation 29 different outcomes of downstream firms. Column 6 presents a regression at the firm-product-time-event level using unit value of inputs as the outcome. In all cases, we restrict the post period to cover up to three months after each flood. Standard errors are clustered at the district level. The top panel presents results for the standardized exposure. The middle panel presents the results for a binary treatment. For supplier sales the binary treatment is whether the supplier was affected by the flood or not. In columns 2-5 the binary treatment is whether the buyer exposure is more than 10 % of purchases. In column 6 the binary treatment is whether the buyer-product exposure is above 10% of purchases. We present the local projections estimates in the bottom panel, where we compute the difference between the post and pre-treatment coefficients. We calculate standard errors for the difference using a bootstrap with 100 repetitions.

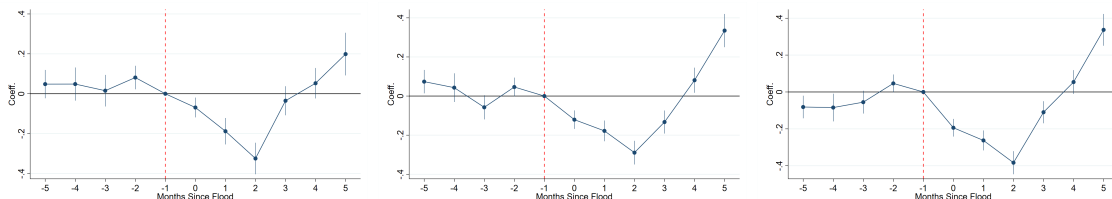
³⁵As our data does not include sales made directly to consumers, we need to impose some additional restrictions to ensure that we focus on firms that consistently sell to other firms. We restrict the sales sample to firms that are observed selling something to other firms every month for the last nine months prior to the flood. We also restrict the sample to be the same as the purchases sample, so we consider the log of 1+sales in cases where the firm is not observed selling anything that period.

Products and input prices. An advantage of our version of the firm-to-firm trade data is that it has detailed product codes and unit values. This allows us to examine product-specific trades and changes in prices as a result of upstream suppliers being exposed to a shock. We first transform the data to the buyer-by-product-by-time level. Our specification is similar to Equation 29, but with a product dimension that allows us to include event-time, industry-district-product-time, and firm-by-product fixed effects, along with controls for pre-period firm-by-product sales interacted with time indicators. In column 6 of Table B4, we study the evolution of product-specific prices for transactions that occur around the flood. While noisier, results are suggestive of a slight increase in price levels three months after the flood when using either the binary treatment or the local projections specification.

New advances in two-way fixed effects methods. Recent econometric advancements in two-way fixed effects methods point out that staggered treatment can lead to the negative weighting of certain disaggregated treatment effects (Goodman-Bacon, 2021). New methods developed by Borusyak et al. (2024); Callaway and Sant’Anna (2020) provide consistent and interpretable estimates. Yet, our setting offers some further challenges. Our “treatment” turns “off” and “on” and perhaps “on” again, and our specifications control for various time-varying covariates, and a wide variety of other fixed effects, making some of these new advances challenging to apply in our setting. A new Local Projections Difference-in-Differences (LP-DID) estimator developed by Dube et al. (2023) allows us to recover interpretable estimates in a flexible and efficient manner.

We present the results from this LP-DID estimator in the bottom panel of Table B4, which show similar patterns. We further implement the LP-DID for the event study analysis as well. In Figure B3b, we once again reproduce the same pattern as before: downstream purchases fall for the first few months, and thereafter recover by month 4. The results from the LP-DID method qualitatively resemble our main results for all other outcomes as well. Figure B3a shows the sales of affected suppliers, and Figure B3c contrasts existing vs. new suppliers. These patterns once again show that sales of affected suppliers fall, and that purchases from buyers decrease from both new and existing suppliers.

Figure B3: LP-DID Event Studies



(a) Sales of affected suppliers (b) Downstream purchases (c) Existing vs new suppliers

Note. Event studies using the Local Projections Difference-in-Differences (LP-DID) approach, discussed in Dube et al. (2023). Figure B3a includes event-time, industry-district-product-time, and firm-by-product fixed effects, and controls for pre-period firm-by-product sales interacted with time indicators. Figure B3b and B3c include firm, time, event-time, and industry-district-real time fixed effects, and demand controls and log pre-period purchases-time controls. Standard errors clustered at the district level.

B.2 Dataset construction

In Section 4.2, we use multiple sources to correlate our model implied probabilities with observables related to supply chain disruption risk. We consider five climate-related measures: rainfall, coastal flooding, riverine flooding, temperature, and drought conditions. Our climate data is available for grid areas that are much more detailed than our 271 regions. We use shape files to overlay our regions to the available maps and calculate the average measure of the climate variables within each of our regions. Coastal and riverine flooding are taken from the World Resources Institute’s Aqueduct Floods Hazard Map. Historical flooding is defined as present-day meters of flooded area. Projected flooding is the 2050 expected meters of increase in flooded areas. We use 10-year floods and the RCP 4.5 as our baseline projection.

Historical and projected temperature and drought data is taken from the IPCC WG1 Interactive Atlas. Historical temperatures are the average daily degrees centigrade in 2005 (the latest year available for historical data). Droughts are measured with the SPI index based on precipitation anomalies over the last 6 months. A lower SPI corresponds to more severe drought conditions. Both temperature and SPI are observed monthly, and we take the average across 12 months to get a value for the gridcell in 2005. Projected data for 2050 is calculated assuming a risk scenario of RCP 4.5 and using a risk model of NOAA global circulation model and the Swedish Meteorological and Hydrological Institute’s local circulation model.

Daily rainfall data is taken from the India Meteorological Department and mea-

sured in millimeters. We take the average across all days in 2019 for each district. For predicted rainfall, we first extract the average historical (measured in 2005) and predicted 2050 rainfall from the IPCC WG1 Interactive Atlas, using the same settings as for temperature. We then compute the change for each district between 2005 and 2050, and apply the implied yearly change to update the 2019 values to 2050.

The non-climate variables mostly come from the Socioeconomic High-resolution Rural-Urban Geographic Platform for India (SHRUG). Elevation is defined as the average elevation in meters of each district while terrain ruggedness is the Terrain Ruggedness Index expressing elevation differences between adjacent pixels. The night-lights luminosity index aims to capture economic activity by detailed regions. Finally, court congestion is taken from the Development Data Lab and measures the average delay in days for the courts in each district.

C Theory Appendix

C.1 Equilibrium Definition

An equilibrium of this economy is a set of state-contingent consumption, $\{q_i(\omega, \boldsymbol{\chi})\}_{\boldsymbol{\chi} \in G(\chi)}$, and final-good labor demand plans, $\{\ell_i^G(\omega, \boldsymbol{\chi})\}_{\boldsymbol{\chi} \in G(\chi)}$, intermediate goods producers labor demands, $\{\ell_i^I\}$, an allocation of input orders, $\{M_{ji}\}_{j \in I, i \in I}$, and a vector of prices and wages, $\{w_i(\boldsymbol{\chi}), p_i^G(\omega, \boldsymbol{\chi}), \mathbb{P}_i(\boldsymbol{\chi}), p_i^I\}_{i \in I, \boldsymbol{\chi} \in G(\chi)}$ such that:

1. Given prices and wages, the representative consumer of each location maximizes its utility.
2. Given prices and wages, firms in each location maximize expected profits.
3. Labor and goods Markets clear state by state

$$\int_{\omega \in [0,1]} \ell_i^G(\omega, \boldsymbol{\chi}) + \ell_i^I = L_i \quad \forall i \in I, \boldsymbol{\chi} \in G(\chi)$$

$$q_i(\omega, \boldsymbol{\chi}) = \phi_i (\ell_i^G(\omega, \boldsymbol{\chi}))^\beta \left(\sum_{j=1} \chi_j M_{ji} \right)^{1-\beta} \quad \forall \omega \in [0, 1], i \in I, \boldsymbol{\chi} \in G(\chi)$$

$$\sum_j \tau_{ij} M_{ij} = z_i \ell_i^I \quad \forall i$$

4. Trade is balanced state by state

$$\sum_j p_j^M \tau_{ji} M_{ji} = \sum_j p_i^M \tau_{ij} M_{ij} \quad \forall i \in I, \boldsymbol{\chi} \in G(\chi)$$

C.2 Proofs

Proposition 1: Proof. Since the cost of materials is linear in M_{ji} and the constraints are conventional (linear) non-negativity constraints, it suffices to show that the expected operating profits function $\mathbb{E}_\chi(\pi(\mathbf{M}_i; \boldsymbol{\chi}))$ is concave in the vector \mathbf{M}_i . The expectation operator preserves the concavity of $\pi(\mathbf{M}_i; \boldsymbol{\chi})$ which is the only thing required to prove. The concavity of ex-post profits, $\pi(\mathbf{M}_i; \boldsymbol{\chi})$, follows from the fact that $\frac{(1-\beta)(\sigma-1)}{\beta+\sigma(1-\beta)} < 1$. ■

Lemmas 1 and 2: Proof. Conditional on some state of the world, $\boldsymbol{\chi}$, ex-post aggregate profits are given by,

$$\int_{\omega \in [0,1]} \pi_i(\omega; \boldsymbol{\chi}) d\omega = \int_{\omega \in [0,1]} \left(p_i(\omega; \boldsymbol{\chi}) q_i(\omega; \boldsymbol{\chi}) - w_i(\boldsymbol{\chi}) \ell_i(\omega; \boldsymbol{\chi}) - \sum_j p_{ji}^M(\boldsymbol{\chi}) M_{ji}(\omega) \right) d\omega.$$

Using the assumption of a unit mass of homogenous firms in a region, ex-post aggregate profits are then

$$\pi_i(\boldsymbol{\chi}) = p_i(\boldsymbol{\chi}) q_i(\boldsymbol{\chi}) - w_i(\boldsymbol{\chi}) \ell_i(\boldsymbol{\chi}) - \sum_j p_{ji}^M(\boldsymbol{\chi}) M_{ji}.$$

where $p_i(\boldsymbol{\chi}) q_i(\boldsymbol{\chi})$ corresponds to aggregate revenues from the final goods sector, $w_i(\boldsymbol{\chi}) \ell_i(\boldsymbol{\chi})$ are payments to labor by final goods producers, and $\sum_j p_{ji}^M(\boldsymbol{\chi}) M_{ji}$ is total expenditure on intermediate inputs.

As final goods firms are monopolistically competitive and the final goods aggregator is CES, standard algebra shows that revenues minus labor costs are a constant fraction of aggregate income:

$$p_i(\boldsymbol{\chi}) q_i(\boldsymbol{\chi}) - w_i(\boldsymbol{\chi}) \ell_i(\boldsymbol{\chi}) = \frac{\beta + \sigma(1 - \beta)}{\sigma} Y_i(\boldsymbol{\chi}).$$

From goods market clearing and trade balance, it is easy to show that aggregate income is equal to the aggregate revenues of the final goods producers, $Y_i(\omega) = p_i(\boldsymbol{\chi}) q_i(\boldsymbol{\chi})$. Plugging this expression in the equation above, we get an aggregate

labor demand equation as a function of wages and aggregate income,

$$\ell_i(\boldsymbol{\chi}) = \frac{\beta(\sigma - 1)}{\sigma} \frac{Y_i(\boldsymbol{\chi})}{w_i(\boldsymbol{\chi})}.$$

Turning to expenditure in intermediates inputs, multiplying the first order conditions defined in Equation 12 by M_{ji} , and adding up across origins j , we obtain:

$$\mathbb{E}_{\boldsymbol{\chi}} \left[\lambda_i(\boldsymbol{\chi}) \left(\Theta_i(\boldsymbol{\chi}) \left[\sum_{j \in I} \chi_j M_{ji} \right]^{\frac{(1-\beta)(\sigma-1)}{\beta+\sigma(1-\beta)}} - \sum p_{ji}^M(\boldsymbol{\chi}) M_{ji} \right) \right] = 0.$$

We can then plug the expression for $\Theta_i(\boldsymbol{\chi})$ and for the stochastic discount factor $\lambda_i(\boldsymbol{\chi}) = \frac{1}{Y_i(\boldsymbol{\chi})}$ to simplify the expression above as:

$$\mathbb{E}_{\boldsymbol{\chi}} \left[\frac{1}{Y_i(\boldsymbol{\chi})} \left(\frac{(1-\beta)(\sigma-1)}{\sigma} Y_i(\boldsymbol{\chi}) - \sum p_{ji}^M(\boldsymbol{\chi}) M_{ji} \right) \right] = 0.$$

Trade balance and zero profits for intermediate goods producers imply that $\sum p_{ji}^M(\boldsymbol{\chi}) M_{ji} = \sum p_{ij}^M(\boldsymbol{\chi}) M_{ij} = w_i(\boldsymbol{\chi}) \ell_i^I$. Thus,

$$\mathbb{E}_{\boldsymbol{\chi}} \left[\frac{(1-\beta)(\sigma-1)}{\sigma} - \frac{w_i(\boldsymbol{\chi})}{Y_i(\boldsymbol{\chi})} \ell_i^I \right] = 0.$$

Imposing labor market clearing, it must be that $\ell_i(\boldsymbol{\chi}) + \ell_i^I = L_i$ for all states of the world. Jointly, with the aggregate demand equation, it follows that

$$L_i - \ell_i^I = \frac{\beta(\sigma-1)}{\sigma} \frac{Y_i(\boldsymbol{\chi})}{w_i(\boldsymbol{\chi})},$$

which in turn, implies that

$$\begin{aligned} \mathbb{E}_{\boldsymbol{\chi}} \left[\frac{(1-\beta)(\sigma-1)}{\sigma} - \frac{\beta(\sigma-1)}{\sigma} \frac{\ell_i^I}{L_i - \ell_i^I} \right] = 0 &\implies \ell_i^I = (1-\beta)L_i \\ \implies \ell_i(\boldsymbol{\chi}) = \beta L_i &\quad \forall i \in I, \boldsymbol{\chi} \in G(\boldsymbol{\chi}). \end{aligned}$$

This means that equilibrium aggregate profits are equal to

$$\begin{aligned}
\pi_i(\boldsymbol{\chi}) &= p_i(\boldsymbol{\chi})q_i(\boldsymbol{\chi}) - w_i(\boldsymbol{\chi})\ell_i(\boldsymbol{\chi}) - \sum_j p_{ji}^M(\boldsymbol{\chi})M_{ji} \\
&= \frac{\beta + \sigma(1 - \beta)}{\sigma} \frac{\sigma}{\beta(\sigma - 1)} w_i(\boldsymbol{\chi})\ell_i(\boldsymbol{\chi}) - w_i(\boldsymbol{\chi})(1 - \beta)L_i \\
&= w_i(\boldsymbol{\chi})L_i \left[\frac{\beta + \sigma(1 - \beta)}{\sigma - 1} - (1 - \beta) \right] \\
&= \frac{w_i(\boldsymbol{\chi})L_i}{\sigma - 1}.
\end{aligned}$$

Finally, from the budget constraint, $Y_i(\boldsymbol{\chi}) = w_i(\boldsymbol{\chi})L_i + \pi_i(\boldsymbol{\chi})$. Combining these expressions, we can show that

$$Y_i(\boldsymbol{\chi}) = \frac{\sigma}{\sigma - 1} w_i(\boldsymbol{\chi})L_i.$$

■

Lemma 3: Proof. Let labor in region 1 be the numeraire. We prove that wages in each location i , w_i , are deterministic by showing that labor market clearing must occur at the time of producing intermediates.

By backward induction, after intermediate inputs have been produced, final goods producers in each region face an inelastic residual labor supply equal to \bar{L}_i . Aggregate labor demand in each region is given by,

$$L_i^D(\boldsymbol{\chi}) = \left[\frac{Y_i(\boldsymbol{\chi})}{\phi_i \left(\sum_{j \in I} \chi_j M_{ij} \right)^{1-\beta} p_i(\boldsymbol{\chi})} \right]^{\frac{1}{\beta}},$$

where final goods' prices can be written as

$$p_i(\boldsymbol{\chi}) = \left[\frac{\beta(\sigma - 1)}{\sigma} \right]^{-\beta} \phi_i^{-1} \left(\sum_{j \in I} \chi_j M_{ij} \right)^{-(1-\beta)} w_i(\boldsymbol{\chi})^\beta Y_i(\boldsymbol{\chi})^{1-\beta}.$$

If we plug the expression for prices, in the aggregate labor demand equation, and simplify we get that,

$$L_i^D(\boldsymbol{\chi}) = \beta L_i$$

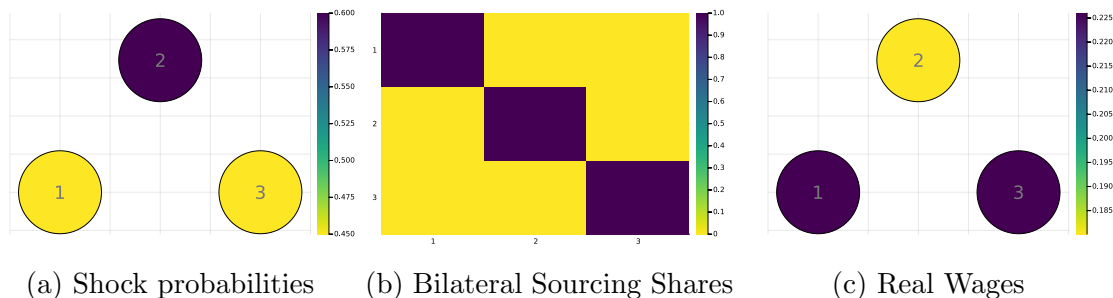
Crucially, aggregate labor demand by final goods producers does not depend on the realization of the shocks, $\boldsymbol{\chi}$. However to clear the labor market in each location the

wage rate needs to be such that the residual labor supply that final goods' producers face, \bar{L}_i , is equal to their inelastic labor demand. The wage rate is set ex-ante when intermediate good production takes place and is independent of the realization of the shocks. As a corollary, this implies that the wage rate, $w_i(\boldsymbol{\chi})$, aggregate profits $\pi_i(\boldsymbol{\chi})$ and aggregate income $Y_i(\boldsymbol{\chi})$ are all deterministic. ■

C.3 Additional Results: Comparative Statics

Heterogeneous risk and autarky. We maintain the scenario in Section 3.4 but raise trade costs to infinity, shutting down inter-regional input sourcing. Figure C1 illustrates that while the probabilities of shocks remain the same as the heterogeneous risk with trade case above, bilateral sourcing mimics a no-risk case. However, the impact on expected real wages is very different. The riskiest region sees the lowest expected real wages, while the safest regions see the highest expected real wages, as they have the lowest expected prices due to the lowest shock probabilities and fully domestic sourcing.

Figure C1: Scenario with heterogeneous risk and infinite trade costs

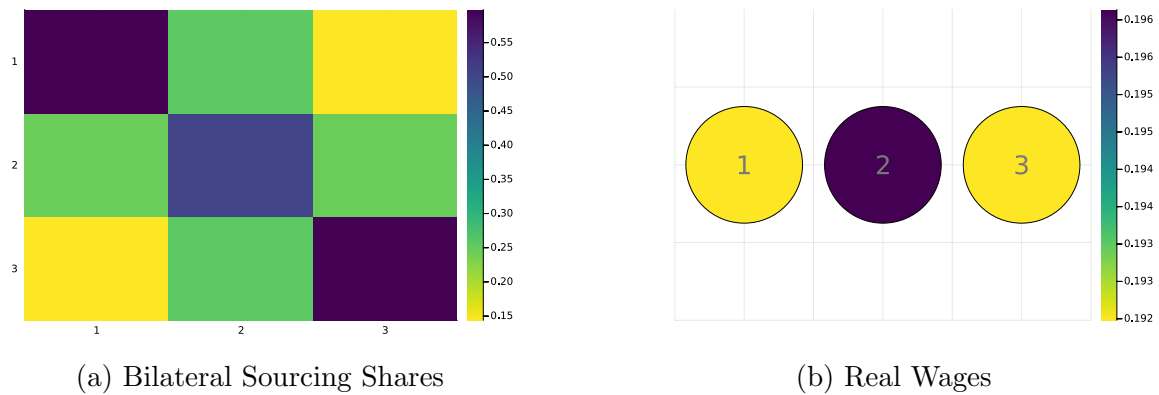


Note. This figure presents the case where trade costs are set to infinity. The figure in the left panel show the probability that each region is hit by a shock, as well as a visual representation of the geographical location of regions in space. The figure in the middle panel consist of a 3x3 input-output matrix where the buying regions are in the vertical axis and the supplying regions are in the horizontal axis. Each line represents the share of inputs purchased by a buying regions from each supplying region. The right panel presents the expected real wages for each region. The scales are shown to the right of each figure.

Homogeneous risk, heterogeneous distance. Figure C2 illustrates the bilateral sourcing shares when the risk of shocks in each region is $\rho = 0.5$. Firms now face a trade-off: as shocks are independent across regions, they can reduce the probability of input disruptions by sourcing from multiple regions. On the other hand, sourcing from other regions is costly, given trade costs. As a result, firms still largely source inputs

from their own regions, but also diversify by sourcing some inputs from geographically closer regions where trade costs are lower. The figure illustrates that this higher demand for inputs from more central regions in equilibrium results in higher expected real wages in these regions. These more central regions also diversify their risk the most by participating in interregional sourcing. Note that the expected price index in more central regions is, therefore, lower in equilibrium, as firms from these regions pay less in trade costs for inputs and better diversify risk.

Figure C2: Scenario with homogeneous risk, heterogeneous distance



Note. The figures in the left panel consist of a 3x3 input-output matrix where the buying regions are on the vertical axis, and the supplying regions are on the horizontal axis. Each line represents the share of inputs purchased by a buying region from each supplying region (column). The figures in the right panel presents the real wages for each region, as well as a visual representation of the geographical location of regions in space. The regions are in a straight line, such that the regions have different distances between each other. The scales are shown to the right of each figure.

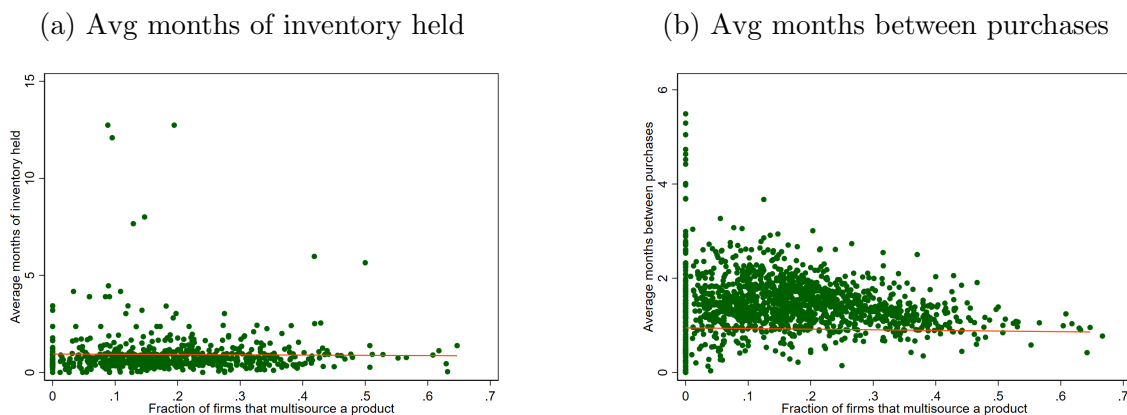
C.4 Inventories - descriptive evidence

In Section B.1.2, we argue that our model without inventories estimates that a 1% decrease in purchases implies a 0.47% decline in sales, which is close to the empirically estimated 0.43% decline in sales. This suggests that while firms might use strategies other than multi-sourcing to protect themselves from shocks, we can approximate the overall sales impact without explicitly incorporating other channels.

We investigate how important inventory holdings are in India. While we do not observe inventories directly in our data, we compute measures at the product level using two alternative approaches. First, we use the 2014-5 Annual Survey of Industries (ASI) to compute, for each HSN-4 product, the average months of inventory held by

firms. We divide the closing value of finished goods by the average monthly sales to measure average inventory/sales. Second, we use our transaction data to compute the average gap in terms of months between two consecutive purchases of each HSN-4 product. Products that are purchased on average with larger gaps will have more accumulated inventories than those with more frequent purchases.

Figure C3: Product-level inventories and multisourcing, by product



Note. In both figures, the horizontal axis plots the share of firms that we observe sourcing a product from at least two suppliers during 2019. In the left panel, the vertical axis measures the average months of inventories held for each product, as computed from the ASI. The vertical axis in the right panel, computes for each product the average number of months between consecutive purchases as measured from our transaction data.

A first thing to note is that the levels of inventories for most products in the data are quite low. According to the average months of inventory held from the ASI, the mean across products is 0.91 months of inventory. The 75th percentile is 0.96, which reinforces that for most product, firms hold less than one months of inventory. In Figure C3, we correlate both measures with the fraction of firms that multisource a given HSN-4 product computed from our transaction data. As shown in the figure, both measures show that there is no correlation between how much a product is multi-sourced and the level of inventory holdings. While inventories might be a relevant strategy for some products, they don't seem to be substitutes or complements to multi-sourcing for our firms.

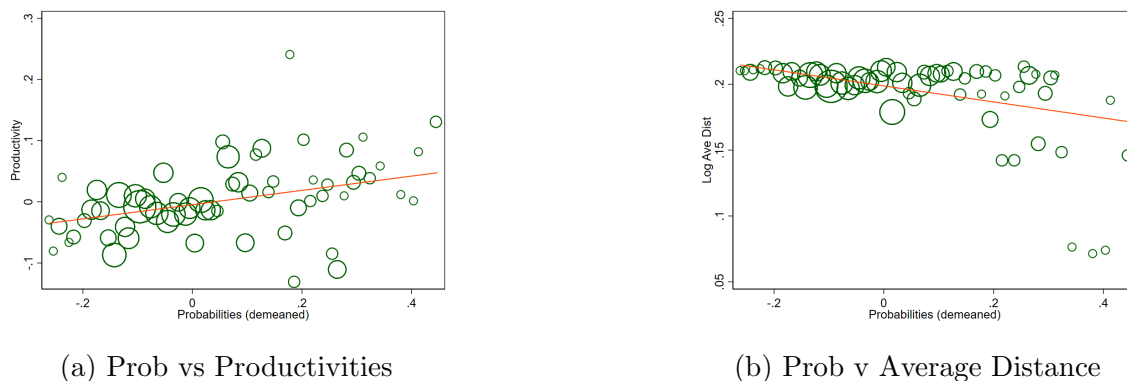
D Quantitative Appendix

Table D1: Calibrated moments

Parameter	Source
L_i : Labor endowments	Annual Survey of Industries (ASI), 2019-20
ϕ_i : Region productivities	Ackerberg et al. (2015) estimation (ASI, 2004-2007)
τ_{ij} : Iceberg trade costs	Regression of within firm-product price on distance between buyer and seller (Transaction data)
ρ_i : Shock probabilities	Minimum distance estimator using sourcing shares across districts (Transaction data)
χ_i : Shock parameter	Match drop in buyer purchases from event study (Transaction data)
β : Labor share	0.19: Ackerberg et al. (2015) estimation (ASI, 2004-2007)
σ : Demand elasticity	2: Based on Boehm et al. (2023)

D.1 Model Probabilities - Additional Analysis

Figure D1: Model probabilities, Productivities and Distance

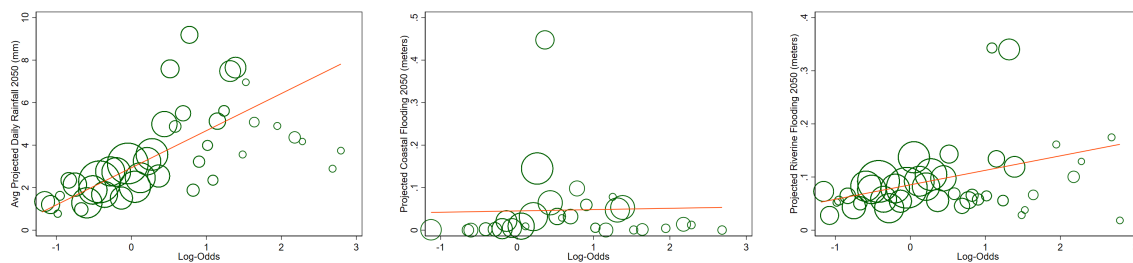


Note. In this figure, we plot the estimated probabilities against some observables. In the left panel, we correlate the probabilities with $\text{Log}(\text{Productivities})$. In the right panel, we correlate the probabilities with the average distance to the state of our study.

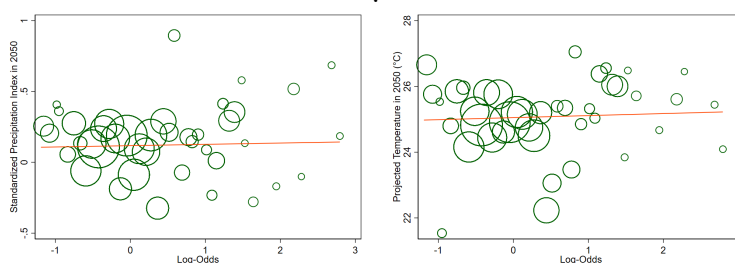
E Alternative Models

In this appendix we consider two alternative models. First, we model risk probabilities using observables in Appendix E.1. Second we consider a CES aggregator of inputs in Appendix E.2

Figure D2: Model Probabilities and Projected Observables



(a) Average Daily Rainfall - Projected 2050 (b) Coastal Flooding - Projected 2050 (c) Riverine Flooding - Projected 2050



(d) Standardized Precipitation Index - Projected 2050 (e) Average Temperature - Projected 2050

Note. In this figure, we plot the estimated probabilities against 2050 projections for climate observables. In Figures D2a and D2e, we correlate the rainfall and temperature projections for year 2050 with the recovered probabilities. Figures D2b use the projected coastal flooding, while Figures D2c correlate the probabilities with projected riverine flooding, respectively. A more detailed definition of each of the variables can be found in Appendix D.1.

Figure D3: Model Probabilities, Price Indices and Wages



(a) Price Index

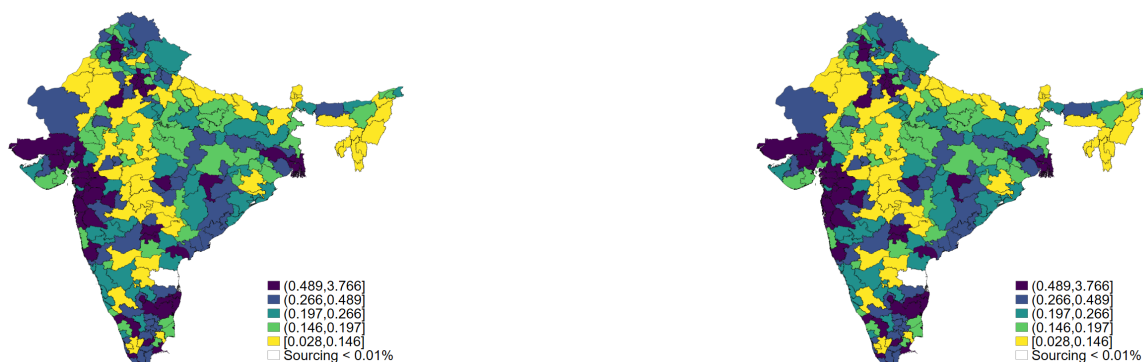
(b) Real Wage Variance

Note. In this figure, we plot the model-derived price index (left panel) and real wage variance (right panel) against the estimated disruption probabilities.

E.1 Projecting Probabilities on Observables

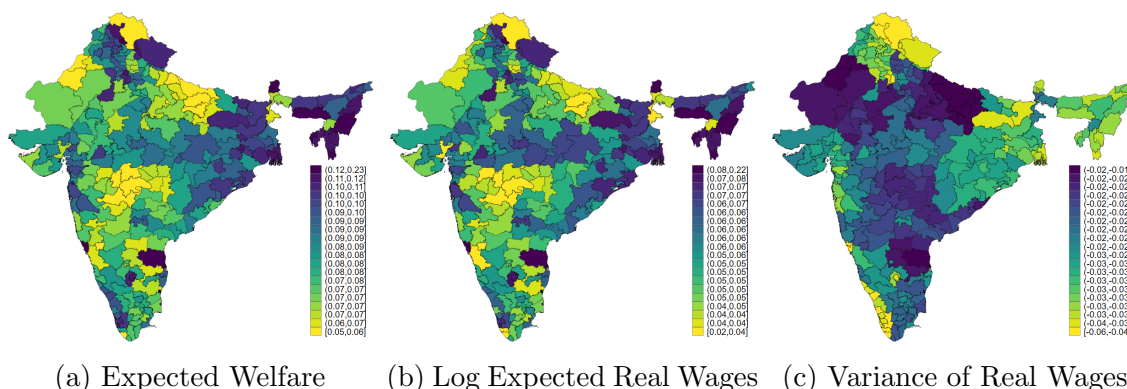
In this model, we describe an alternative estimation strategy for the disruption probabilities, ρ_i . Instead of computing one parameter per region, we parameterize the

Figure D4: Ahmadabad sourcing, Kolkata sourcing – Free Trade



Note. In this figure, we plot model sourcing shares with Free Trade for Ahmadabad district (left panel) and Kolkata district (right panel).

Figure D5: Quantitative results - Δ in Free Trade



(a) Expected Welfare (b) Log Expected Real Wages (c) Variance of Real Wages

Note. This figure shows welfare (Panel a), expected real wages (Panel b) and their variance (Panel c) for the counterfactual of free trade. The figures show the percentage changes in these variables under the free trade counterfactual relative to the baseline scenario.

vector $\{\rho_i\}_{i=1}^I$ on a vector of observable characteristics, Z_i . This vector Z_i includes a constant term, average court delays, ruggedness, elevation, night lights, average rainfall, average coastal flooding, average riverine flooding, and average temperature. We include all of these variables in logs, and we add a dummy for the case in which historical coastal flooding is positive, to allow the function to allow the function to flexibly estimate the asymptotic behavior of the log at 0. Then, we assume that these probabilities have the following functional form,

$$\rho_i = \frac{e^{Z_i' \gamma}}{1 + e^{Z_i' \gamma}},$$

where γ is the vector of parameters that we estimate by minimizing the gap between model-implied and the observed average sourcing shares in the data.

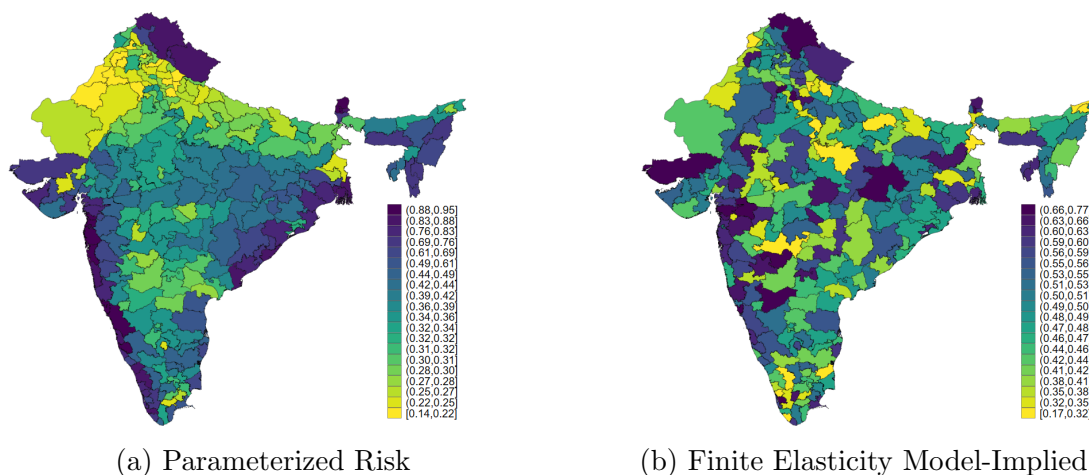
In Table E2, we present the estimates of the vector γ . The resulting probabilities from this approach are shown in Panel (c) of Figure 5.

This estimation approach requires estimating fewer parameters than our baseline, but necessitates that we take a stance on the sources of district-level risk. While the estimation approaches are independent of each other, the estimated coefficients for rainfall, flooding and temperature are positive, consistent with the baseline. Night-lights have a zero coefficient, also consistent with the baseline. In contrast to the baseline, however, courts also contribute positively to risk under this approach.

Table E2: Estimates of the Model for the Probabilities

	Constant	Courts	Ruggedness	Elevation	Night Lights	Rainfall	Coastal Flooding	Coastal Dummy	Riverine Flooding	Temperature
γ	-1.20	0.01	1.18	0.08	0.00	0.19	0.27	0.08	0.07	0.82

Figure E6: Estimated Disruption Probabilities for Alternative Models



Note. We plot the model-implied district-level disruption probabilities for the alternative models. The left panel plots the district-level disruption probabilities implied by the parameterized approach outlined in the text. The right panel shows the district-level disruption probabilities obtained by following the same approach as in the baseline model, but allowing a finite elasticity of substitution across inputs of different origins. The scales are shown to the right of each figure.

E.2 A Model with Finite Elasticity Across Inputs

In this appendix, we develop a model in which we relax the assumption of perfect substitution of inputs across different regions by allowing for a finite elasticity of substitution, akin to an Armington model. Firms will have two incentives to source input varieties from different regions. The first one is the diversification motive, which is the main focus of this paper. The second incentive corresponds to love-for-variety. The only modification to the model in Section 3 is to allow for imperfect substitution in the aggregator of inputs in Equation 6. Thus, the expression becomes:

$$x_i(\omega) = \left(\sum_{j \in I} x_j^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}}.$$

Since this assumption is just changing the way that the received input units are aggregated, the *ex-post* problem of the firm remains unchanged. Profits as a function of the total number of inputs the firm has are:

$$\pi_i(\mathbf{M}_i; \boldsymbol{\chi}) = \kappa w_i^{\frac{\beta(1-\sigma)}{\beta+\sigma(1-\beta)}} \left[[Y_i \mathbb{P}_i^{\sigma-1}] \phi_i^{\sigma-1} \left(\left(\sum_{j \in I} [\chi_j M_{ji}]^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}} \right)^{(1-\beta)(\sigma-1)} \right]^{\frac{1}{\beta+\sigma(1-\beta)}},$$

where $\kappa = \left[\frac{\sigma(1-\beta)+\beta}{\beta(\sigma-1)} \right] \left[\frac{\beta(\sigma-1)}{\sigma} \right]^{\frac{\sigma}{\beta+\sigma(1-\beta)}}$. The sourcing problem of the firm is to choose M_{ij} to maximize expected profits minus order costs

$$\max_{M_{ij} \geq 0} \mathbb{E}_{\boldsymbol{\chi}} \left(\kappa w_i^{\frac{\beta(1-\sigma)}{\beta+\sigma(1-\beta)}} \left[[Y_i \mathbb{P}_i^{\sigma-1}] \phi_i^{\sigma-1} \left(\left(\sum_{j \in I} [\chi_j M_{ji}]^{\frac{\varepsilon-1}{\varepsilon}} \right)^{\frac{\varepsilon}{\varepsilon-1}} \right)^{(1-\beta)(\sigma-1)} \right]^{\frac{1}{\beta+\sigma(1-\beta)}} - \sum_{j \in I} p_j^M M_{ji}, \right) \quad (30)$$

with first-order condition,

$$\mathbb{E}_{\boldsymbol{\chi}} \left(\chi_j \Theta_i \left[\sum_{j \in I} (\chi_j M_{ji})^{\frac{\varepsilon-1}{\varepsilon}} \right]^{\frac{-\varepsilon+\beta+\sigma(1-\beta)}{\beta+\sigma(1-\beta)}} (\chi_j M_{ji})^{-\frac{1}{\varepsilon}} \right) \leq p_j^I.$$

In this particular model due to an Inada condition, the solution will be interior, and is implicitly given by (after plugging in the GE components):

$$M_{ji} = (1 - \beta)^\varepsilon (w_i L_i)^\varepsilon \frac{(p_j^I)^{-\varepsilon}}{\mathbb{E}_\chi \left(\chi_j^{\frac{\varepsilon-1}{\varepsilon}} \left[\sum_{j \in I} (\chi_j M_{ji})^{\frac{\varepsilon-1}{\varepsilon}} \right]^{-1} \right)^{-\varepsilon}}.$$

where

$$\Theta_i = (1 - \beta) w_i L_i \left(\sum_{j \in I} [\chi_j M_{ji}]^{\frac{\varepsilon-1}{\varepsilon}} \right)^{-\frac{\varepsilon}{\varepsilon-1} \frac{(1-\beta)(\sigma-1)}{\beta+\sigma(1-\beta)}}$$

Notice that we cannot derive a closed-form solution for this expression; we can only define it implicitly, and solve for the demand of inputs numerically.

Proposition 2 *The ex-ante profit function described in Equation 30 is concave in orders of inputs M_{ji} .*

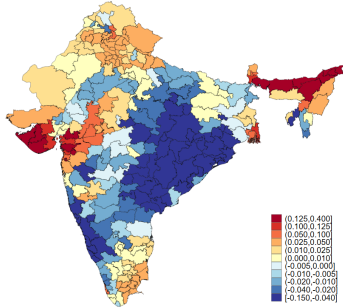
Proof. As the cost of materials is linear in M_{ij} and constraints are conventional (linear) non-negativity constraints, it suffices to show that the expected profits function $\mathbb{E}_\chi (\pi(\mathbf{M}; \boldsymbol{\chi}))$ is concave in \mathbf{M} . The expectation operator preserves the concavity of $\pi(\mathbf{M}; \boldsymbol{\chi})$ which is the only thing required to prove. Concavity of the CES aggregator follows from the fact that it is a quasi-concave function homogeneous of degree 1. The concavity of ex-post profits, $\pi(\mathbf{M}; \boldsymbol{\chi})$, follows from the parametric restriction, $\frac{(1-\beta)(\sigma-1)}{\beta+\sigma(1-\beta)} < 1$, as the composition of concave functions is concave. ■

E.3 Quantitative Implications

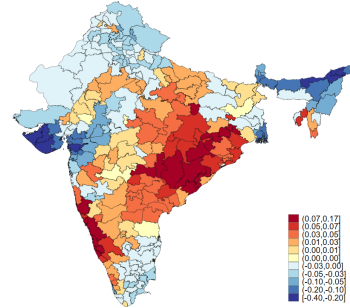
Table E3 summarizes the baseline and counterfactuals in the two alternative models. The insights are similar to the baseline model. In both models, autarky is welfare decreasing, though there is spatial heterogeneity. In the CES model, autarky decreases welfare by two orders of magnitude more as autarky additionally implies losses from variety as only own-region inputs can be used to produce. Free trade is welfare improving in both models. Interestingly, the implications of climate risk changing are similar in both models, despite their independent estimation and varied structure. On average welfare decreases by 2% in the climate counterfactual in both models (2.01% in the baseline). The fraction of districts with real wage declines is larger in the CES model, at 86.35%, than in the projected probabilities model, at 54.98%.

Figure E7: Counterfactuals: Climate Risk Increase – Alternative Models

Parameterized Risk Profile

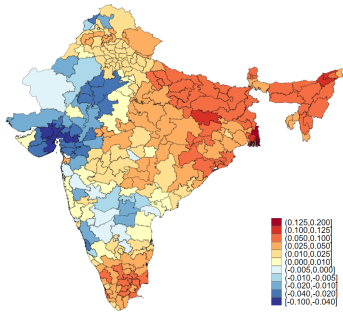


(a) Δ in Probabilities

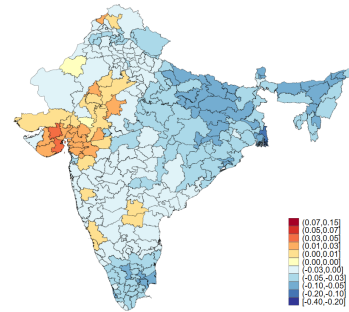


(b) Δ in Welfare

Finite Elasticity of Substitution Across Inputs



(c) Δ in Probabilities



(d) Δ in Welfare

Note. We plot the change in probabilities of climate risk (Panel A), and the change in welfare (Panel B) for the model with parameterized risk. In Panel C and Panel D, we plot the change in probabilities of climate, and the change in welfare for the model with a finite elasticity of substitution.

Table E3: Model Counterfactuals: Summary – Alternative Models

Counterfactual	Δ in Welfare		Δ in log Expected Real Wages		Δ in Real Wage Volatility		% districts Real wage declines
	Avg. change	Range	Avg. change	Range	Avg. change	Range	
<u>Parameterized Risk</u>							
<u>Baseline risk</u>							
Autarky	-6.79	3.75	2.54	1.31	7.14	5.97	2.58%
Free Trade	7.40	2.32	4.96	1.57	-2.42	1.32	0.00%
<u>Alternative risk</u>							
Climate change	-2.00	4.68	-2.12	4.67	-0.02	0.11	54.98%
<u>Finite Elasticity of Substitution Across Inputs</u>							
<u>Baseline risk</u>							
Autarky	-198.96	42.01	-186.83	43.34	11.87	3.71	100.00%
Free Trade	15.33	1.08	15.33	1.07	4.25	2.45	0.00%
<u>Alternative risk</u>							
Climate change	-2.00	2.71	-1.96	2.71	0.03	0.00	86.35%

Note. This table shows statistics of the distribution of percentage changes between the baseline scenario with current climate risk and costly trade and other scenarios, weighted by district population. Range refers to the interquartile range.

References

- Ackerberg, D. A., K. Caves, and G. Frazer (2015, November). Identification Properties of Recent Production Function Estimators. *Econometrica* 83(6), 2411–2451.
- Alfaro Ureña, A., M. Fuentes Fuentes, I. Manelici, and J. P. Vasquez (2018). Costa Rican Production Network: Stylized Facts. *Working paper*.
- Boehm, C. E., A. A. Levchenko, and N. Pandalai-Nayar (2023, April). The Long and Short (Run) of Trade Elasticities. *American Economic Review* 113(4), 861–905.
- Borusyak, K., X. Jaravel, and J. Spiess (2024). Revisiting Event Study Designs: Robust and Efficient Estimation. *Review of Economic Studies*.
- Callaway, B. and P. H. C. Sant’Anna (2020). Difference-in-Differences with Multiple Time Periods. *Journal of Econometrics*.
- Dube, A., D. Girardi, O. Jorda, and A. Taylor (2023). A Local Projections Approach to Difference-in-Differences Event Studies. *NBER Working Paper 31184*.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics* 225(2), 254–277.