Beyond Targeted Firms: Supply Chain Spillovers of Environmental Regulation in China

Kai Wu *

This Draft: March 2026

(Click HERE for the most recent draft)

Abstract

This paper studies how regulatory shocks propagate through production networks and who bears the burden. I exploit a natural experiment—the 2017 Xiong'an pollution shutdown campaign in China, which imposed mandatory closures on polluting industrial operations across a newly designated development zone—combined with administrative VAT data covering the universe of firm-to-firm transactions. Targeted firms' purchases fell by 55% and sales by 53%. The shock propagated in both directions: upstream suppliers at mean exposure lost 11.8% of sales; downstream customers saw purchases drop by 9.5%. Effects extend to second-degree trade partners and persist for at least three years. Small firms bear the entire propagation loss while large firms are completely buffered. The mechanism is two-sided complementarity: for small firms, trading partners are gross complements on both the input and output sides, so that losing one partner forces contraction with all others; for large firms, partners are gross substitutes. A general equilibrium model of shock transmission in production networks with scale economies and size-varying elasticities formalizes these patterns, showing that ignoring supply chain spillovers leads to overregulation and that network externalities create novel channels through which emissions taxes dominate quantity controls.

Keywords: Production networks, environmental regulation, shock propagation, firm heterogeneity, two-sided complementarity, optimal policy design.
JEL Codes: D22, H23, L14, L51, Q58

---

[*] Department of Economics, University of Texas at Austin. Email: kaiwu@utexas.edu. I am grateful to Marika Cabral, Jackson Dorsey, Nitya Pandalai-Nayar, and Cody Tuttle for their guidance and support. I thank Saroj Bhattarai, Jaedo Choi, Olivier Coibion, Michael Geruso, Husang Kim, Leigh Linden, Rikuto Onishi, Bradley Shapiro, Ben Zou, and seminar participants at the University of Texas at Austin for helpful comments. This draft is a work in progress and comments are welcome; all errors are my own.

1 Introduction

Regulations impose costs on targeted firms, and a large literature measures these costs by examining direct impacts on the regulated entities—their productivity, output, employment, or profits. But firms are embedded in production networks. When regulation restricts a firm's operations, the consequences do not stop at that firm: upstream suppliers lose demand, downstream customers lose inputs, and these disruptions cascade outward. The spillovers can be large, especially when trading partners are gross complements—so that losing one partner forces contraction with all others—and when forming new relationships is costly. Despite this, evidence on the network spillovers of regulation remains scarce, both in the literature and in policymaking practice, where standard cost-benefit analyses quantify compliance costs at regulated facilities but omit losses propagated through input-output linkages.[1]

To fill this gap, this paper studies how regulatory shocks propagate through production networks, who bears the burden, and what these spillovers imply for optimal policy design. I exploit a unique natural experiment in China: the pollution shutdown campaign associated with the establishment of the Xiong'an New Area in 2017. Three institutional features offer the opportunity for clean identification. First, the campaign was part of a confidential central government decree; unlike routine environmental inspections, which often target firms based on past performance, observable pollution levels, or local economic conditions, the Xiong'an shutdowns were predetermined by geographic boundaries established through a confidential planning process. Local governments and market participants were not notified in advance, eliminating concerns about anticipatory responses or selection based on unobserved firm characteristics. Second, the policy was implemented with unprecedented stringency as part of a top-level national priority. By late 2017, over 10,000 firms in the area had been permanently shut down, creating rich variation that allows for identifying both direct and propagation effects. Third, unlike natural disasters (Barrot and Sauvagnat, 2016; Carvalho et al., 2021; Boehm et al., 2019), which simultaneously damage infrastructure and disrupt transportation, or pandemics (Barrot et al., 2021; Baqaee and Farhi, 2022; Khanna et al., 2022), which bring confounding demand shocks and labor supply constraints, the regulatory shutdown affected only the targeted firms' ability to operate. The broader economic environment remained intact, making it possible to trace the pure transmission of supply and demand shocks through the production network.[2]

---

[1] For example, the U.S. EPA's *Guidelines for Preparing Economic Analyses* (2010, Ch. 8) adopts a partial equilibrium framework in which "indirect costs" consist of price pass-through and deadweight loss in the regulated market. The 2023 revision of OMB Circular A-4, the first update in 20 years, expanded its scope to distributional effects and market power but does not direct agencies to quantify network externalities transmitted through input-output linkages. The 2023 Regulatory Impact Analysis for the Mercury & Air Toxics Standards (MATS) illustrates this gap: it quantifies scrubber installation and fuel-switching costs at affected power plants and health benefits from reduced pollution, but omits upstream supplier losses (e.g., coal miners, equipment manufacturers) and downstream customer disruptions propagated through the production network.

[2] Although the Xiong'an campaign is set in China, the phenomenon it illuminates is not unique to the Chinese context. Command-and-control environmental enforcement—including forced closures of polluting facilities—is common

The empirical analysis draws on administrative value-added tax (VAT) data covering the universe of firm-to-firm transactions in China from 2014 to 2018. The dataset encompasses approximately 16.1 billion transactions among roughly 16 million establishments across all sectors of the economy, and its granularity is central to the analysis. First, it allows me to map each affected firm's supply chain before the policy announcement—who its direct trading partners were, and how those partners connected to the rest of the network. Second, by aggregating transactions within each firm, I can track overall performance over time; by decomposing trade across partners, I can observe how firms reallocate activity within their existing networks and whether they form new connections after a shock. Third, the transaction-level information makes it possible to separate price changes from quantity adjustments, which is valuable for validating the microfoundations of the structural model and estimating its parameters.

I use a difference-in-differences strategy to estimate both the direct effect on the targeted firms and the propagation effects on their trading partners. For the direct effect, I compare Xiong'an industrial firms to a matched control group constructed via propensity score matching on pre-campaign firm characteristics. The specification includes firm fixed effects, industry-province-time fixed effects, and matching-group-time fixed effects, so that the estimates are identified from within-firm variation after absorbing common regional-industry and matching group shocks. Pre-policy coefficients are small and statistically indistinguishable from zero, confirming the parallel trends assumption. Following full implementation in October 2017, purchases and sales of the targeted firms declined by approximately 55% and 53%, respectively.[3] This sharp contraction provides the empirical foundation for investigating the network spillover effects.

For the propagation analysis, I measure a firm's exposure to the shutdown as the share of its pre-policy sales to (for upstream suppliers) or purchases from (for downstream customers) the soon-to-be-shutdown Xiong'an firms. The results show propagation in both directions. Upstream suppliers at sample-mean exposure experienced an 11.8% decline in sales and a 7.7% reduction in purchases; downstream customers saw purchases fall by 9.5% and sales contract by 7.9%. The shock also reaches second-degree partners—suppliers of suppliers and customers of customers—with attenuated per-firm effects but a far larger set of affected firms: the second-degree samples are roughly 40 times

---

worldwide: India's Supreme Court ordered the closure of over 1,300 polluting industries in the Delhi-NCR region; the European Union has legislated mandatory coal phase-outs (Germany by 2038, UK completed 2024); and in the United States, the Clean Air Act and Mercury & Air Toxics Standards have triggered waves of plant shutdowns in non-attainment counties. More broadly, any regulation that reduces a firm's output—whether through taxes, emission caps, or technology mandates—creates the same upstream demand loss and downstream supply disruption through the same network channels.

[3] The decline falls short of 100% for three reasons. First, my treatment group is defined as all industrial firms located in Xiong'an, which casts a wide net and may include "cleaner" firms that were not directly targeted by the campaign. Second, "shutdown" in the context of the campaign refers to the polluting production line, not necessarily the entire firm—some firms retained non-polluting operations, upgraded their technology, or relocated. Third, even firms that were fully shut down may generate residual VAT transactions during the wind-down period, such as liquidating inventory, settling outstanding contracts, or selling off equipment.

(upstream) and 70 times (downstream) as large as the corresponding first-degree samples. Because the VAT data end in 2018, I supplement them with total revenue from annual firm financial statements (available through 2020) to extend the post-policy window. The revenue-based estimates are comparable in magnitude to the VAT-based results for both first- and second-degree partners. These negative effects persist for at least three years with no sign of recovery through the end of the sample period, consistent with sticky production networks in which forming new trade relationships is costly (Huneeus, 2018; Bernard et al., 2019; Elliott et al., 2022).

Heterogeneity analysis reveals that firm size is the dominant driver of how these shocks transmit. Small firms (below-median sales) bear the full burden of the propagation shock, while large firms are completely buffered—the implied net effect on large firms is indistinguishable from zero across all specifications. A back-of-envelope aggregation that accounts for size-varying elasticities, size-varying trade exposure, and the skewness of the sales distribution puts first-degree spillover losses at approximately 84% of the direct loss on Xiong'an industrial firms (46% from upstream, 39% from downstream). Small firms, despite accounting for only 1.2% of total group sales, bear essentially all of this spillover loss—a pattern I term "Network Regressivity."

The mechanism behind this pattern is size-varying complementarity across trade relationships. For small firms, trading partners are gross complements on both the input and output sides: losing one partner causes them to reduce trade with all remaining partners. For large firms, partners are gross substitutes: losing one has no effect on—or even slightly increases—trade with the others. On the input side, this is the standard channel emphasized in the production networks literature (Acemoglu et al., 2012; Carvalho et al., 2021; Baqaee and Farhi, 2022): when inputs are complements, losing one supplier constrains the use of all others. On the output side, the complementarity is, to my knowledge, a novel finding: when a small upstream supplier loses a customer to the shutdown, it also reduces sales to its other, unaffected customers. Standard input complementarity cannot account for this pattern. Instead, it reflects scale economies in the broad sense: a firm that loses a major customer sees its effective scale shrink, raising average costs and eroding productive capacity, which in turn limits its ability to serve remaining customers. The variation in complementarity across firm size can be partially traced to the availability of pre-existing alternative partners, which is strongly correlated with firm size. Large firms tend to have more existing alternatives for each product to reallocate when one partner is lost than small firms. I also investigate whether firms adjust on the extensive margin by forming new relationships. The pace of new connection formation after the shock is slow and falls far short of offsetting the lost trade, consistent with costly network formation (Huneeus, 2018; Bernard et al., 2019; Elliott et al., 2022). This confirms that propagation flows primarily through existing networks and that complementarity among existing partners is what drives the spillover losses.

To formalize these mechanisms and draw out their welfare implications, I develop a general equilibrium model of production networks building on Carvalho et al. (2021), in which competitive firms produce using labor and a CES aggregate of intermediate inputs purchased from other firms, and shocks propagate through input-output linkages. I extend their framework with two features to match my reduced-form findings: scale economies and size-varying elasticities of substitution. The first captures the output complementarity documented in the data. I model scale economies as demand-dependent TFP: when a firm's customers contract, its effective productivity erodes, and this erosion feeds back through the network because the firm's own suppliers now face a less productive buyer. The resulting effective TFP is the solution to a fixed-point problem in which each firm's productivity depends on the productivity of the firms it sells to, generating a cascade of amplification. This demand–productivity linkage is stronger for small firms (whose customer base is concentrated) than for large firms (whose customers are diversified). The second extension allows the elasticity of substitution across intermediate inputs (ξ_i) to vary with firm size, nesting the empirical finding that small firms treat suppliers as gross complements ($\xi_S < 1$) while large firms treat them as gross substitutes ($\xi_L > 1$). Together, these extensions generate network regressivity as a structural result: the ratio of small-to-large Domar weight losses is amplified at each stage.

I estimate five structural parameters—σ (labor vs. intermediates substitution), ξ_0 and ξ_1 (input substitution, base and size differential), φ_0 and φ_1 (scale economies, base and size differential)—by minimum distance, matching model predictions to six within-firm ratio moments constructed from the same samples and difference-in-differences identification as the reduced-form analysis. The identification structure is block-diagonal: the purchase-to-sales ratio identifies σ (with one overidentifying restriction), the affected-to-unaffected expenditure ratio identifies (ξ_0, ξ_1) following Boehm et al. (2019), and upstream firm price changes identify (φ_0, φ_1). Each block can be estimated independently, with no cross-contamination. Because the model-predicted moments are linear in the parameters, estimation reduces to closed-form weighted least squares.

The welfare analysis shows that network structure matters for the design of environmental policy. I evaluate policy counterfactuals using a social welfare function that balances GDP against pollution damages, calibrating the marginal social benefit of pollution reduction from revealed-preference willingness-to-pay estimates in Ito and Zhang (2020). Network effects enter on both sides of the ledger: the social cost of reducing a firm's output includes propagated losses through scale and complementarity channels; the social benefit includes "incidental" emission reductions from contracting supply chain partners. Because regulation targets dirty firms while their upstream suppliers and downstream customers are systematically cleaner, the GDP cascade is not matched by proportional emission reductions—implying that a regulator who ignores network spillovers will overregulate. Production networks also introduce a source of tax advantage distinct from the standard selection-on-abatement-costs mechanism: under two-sided complementarity, the GDP loss function is convex in firm-level shocks, and demand-dependent TFP creates cascade externalities that a uniform quota ignores but a tax avoids by shielding central firms.

Related Literature. This paper contributes to three literatures. First, it contributes to research on shock transmission through production networks. Theoretical work has established that input-output linkages can amplify idiosyncratic shocks into aggregate fluctuations (Acemoglu et al., 2012; Baqaee and Farhi, 2019). On the empirical side, one strand traces propagation from natural disasters such as the Tōhoku earthquake and major hurricanes (Barrot and Sauvagnat, 2016; Carvalho et al., 2021; Boehm et al., 2019; Castro-Vincenzi et al., 2024); another studies policy-induced disruptions, including pandemic lockdowns (Barrot et al., 2021; Baqaee and Farhi, 2022; Khanna et al., 2022), tariffs (Flaaen and Pierce, 2024), and government spending (Barattieri et al., 2023). A parallel literature treats network structure as endogenous and studies the formation and evolution of firm-to-firm linkages (Huneeus, 2018; Bernard et al., 2019, 2022; Elliott et al., 2022; Miyauchi, 2024). I contribute along three dimensions. First, I use administrative firm-to-firm transaction data to study an environmental regulatory shock—a setting that provides cleaner identification than natural disasters (no confounding infrastructure damage) or pandemics (no aggregate demand effects) and carries direct policy implications. Second, beyond documenting propagation in both directions and out to second-degree partners, I show that the burden falls entirely on small firms while large firms are buffered, and trace this asymmetry to the fact that small firms treat trading partners as gross complements on both the input and output sides whereas large firms treat them as substitutes.[4] Third, I extend the general equilibrium framework of shock transmission through input substitution (Baqaee and Farhi, 2019; Carvalho et al., 2021) by incorporating output-side complementarity and size-varying elasticities.

Second, it contributes to the literature on the economic costs of environmental regulation. Traditional studies focus on direct effects on regulated entities (Gray, 1987; Jaffe et al., 1995; Greenstone, 2002; Walker, 2013; He et al., 2020). Recent work has documented two spillover channels: intra-firm reallocation through ownership networks, where regulation on one plant shifts activity to co-owned plants elsewhere (Bartram et al., 2022; Cui and Moschini, 2020; Chen et al., 2025; Curtis et al., 2025); and inter-firm reallocation through market competition, where unregulated competitors expand as regulated firms contract (Fowlie, 2009; Fowlie et al., 2016; Fell and Maniloff, 2018). The channel I identify—spillovers through input–output linkages—is distinct from both. Reallocation effects redistribute production within corporate boundaries or across market competitors; the network spillovers I document are net contractions that propagate through the supply chain. They are neither internalized by regulated firms nor offset by competitors, and my estimates suggest they at least double the conventionally measured cost of regulation.

Third, it contributes to optimal environmental policy design by bringing production networks into the welfare calculation. The standard framework derives optimal Pigouvian taxes or quantity

---

[4] This finding also speaks to the literature on rising industrial concentration (Autor et al., 2020; De Loecker et al., 2020; Kwon et al., 2024; Choi et al., 2024). Most of that work asks why large firms grow, pointing to causes such as technology adoption, globalization, markups, and network-based accumulation of customers and suppliers (Bessen, 2020; Hsieh and Rossi-Hansberg, 2023; Melitz, 2003; Bernard et al., 2022). I document a mechanism on the failure side in the context of production networks: negative shocks act as a selection device that favors large, well-connected firms over small, peripheral ones.

controls by equating marginal abatement cost with marginal social benefit (Pigou, 1920), with a large literature extending this framework along dimensions such as market power, heterogeneous firms, pre-existing tax distortions, stock pollutants, and carbon leakage (Buchanan, 1969; Weitzman, 1974; Fowlie, 2010; Goulder et al., 1999; Goulder and Parry, 2008; Hoel and Karp, 2002; Fell and Maniloff, 2018; Fontagné and Schubert, 2023). To my knowledge, no prior work has incorporated production network structure into the optimal policy problem. I show that ignoring network externalities leads to overregulation—the social marginal cost of abatement is steeper than a partial-equilibrium calculation suggests—and that the network channel introduces novel sources of tax advantage beyond the standard selection mechanism.

The remainder of the paper is organized as follows. Section 2 describes the institutional background and data sources. Section 3 presents the research design. Section 4 provides descriptive statistics. Section 5 presents the main empirical results on direct and propagation effects and conducts robustness checks. Section 6 documents heterogeneity by firm size and quantifies aggregate spillover losses. Section 7 investigates the mechanisms—size-varying complementarity, pre-existing alternatives, and network stickiness. Section 8 develops the structural model and describes the estimation strategy. Section 9 presents the welfare analysis and compares policy instruments. Section 10 concludes.

2 Background and Data

2.1 Background: Xiong'an New Area Plan & Pollution Shutdown Campaign

To support the development of a "green" new city in the Xiong'an New Area, authorities launched an intensive pollution shutdown campaign in April 2017. The campaign targeted local pollution sources with the stated goal of eliminating industrial pollution within the area. Identified firms faced three regulatory paths: permanent shutdown, temporary shutdown with mandated abatement upgrades, or relocation outside the area. In practice, permanent shutdowns predominated. The campaign unfolded in two broad phases: (i) April–September 2017 (inspection and designation), when polluters were identified and assigned regulatory options; and (ii) October 2017 onward (full implementation), when shutdowns were executed and, from 2019, oversight shifted toward preventing pollution resurgence. As summarized in Table 1, by October 2017, over 10,000 firms had been permanently shut, a number that continued to grow through 2018.

This environmental campaign was a core component of the broader Xiong'an New Area plan, which had been announced by the Chinese government on April 1, 2017. The New Area is a state-level new area in Baoding, Hebei Province, roughly 100 km (62 miles) south of Beijing, consisting of three counties—*Xiong, Rongcheng, and Anxin*—and three adjacent townships. Xiong'an is primarily intended to accommodate nonessential functions relocated from Beijing, including state-owned enterprises (SOE), peripheral government agencies, public universities, and research and development (R&D) facilities, and to serve as an economic hub for coordinating development of the Beijing–Tianjin–Hebei (*Jing–Jin–Ji*) region.

Before the announcement, the initiative and preparation for the plan were carried out under a high degree of confidentiality by the central government; local governments were not notified in advance. Site selection emphasized regional centrality and ecological restoration. The area sits near the geometric center of the *Jing–Jin–Ji* triangle and is anchored on *Baiyangdian* (Baiyang Lake), northern China's largest freshwater lake. The plan mandated wetland rehabilitation and strict environmental controls, which motivated the shutdown campaign.

Table 1: Xiong'an pollution shutdown campaign: timeline and statistics

Date	Milestone / action	Stats
May 2017	First public reports of regulatory actions	–
2017-07-08	Firms identified and closures to date	7,248 identified; 3,909 shut
2017-08-03	Targeted month-long enforcement launched	–
2017-10-12	Firms identified and closures to date	11,387 identified; 10,661 permanently shut; 726 temporarily/partially closed for improvements
2018-10-29	Additional firms identified and actions	+1,433 identified; +915 shut; +518 required improvements
End-2018	Authorities announce polluting firms "basically eliminated"; focus shifts to prevention onwards	–

Source: Official announcements and contemporaneous news reports. "Firms" include numerous micro- and self-employed entities, many below the RMB 30,000 average monthly sales threshold, which are not captured in my VAT invoice database.

Figure 1: Map of the Xiong'an New Area

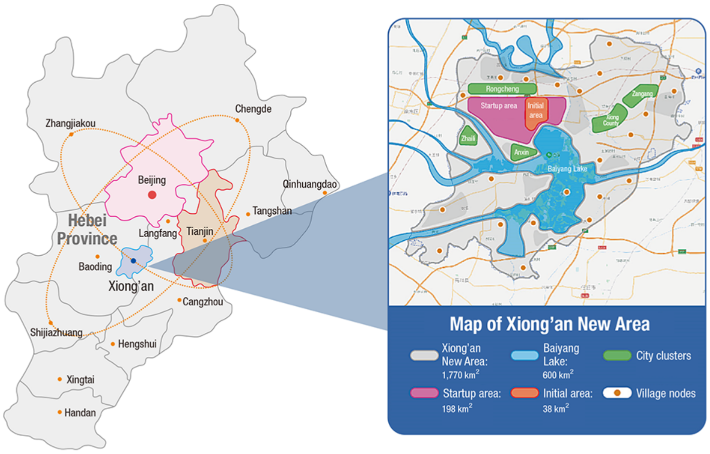

Prepared by Design Office of Hitachi (China) Research & Development Corporation (HCR&D)

Notes: The figure shows the location of the Xiong'an New Area relative to Beijing and Tianjin, and a detailed map of its three constituent counties and the Baiyang Lake. *Source:* Hitachi (China) Research & Development Corporation (HCR&D).

2.2 Data

This study links three administrative sources: (i) the universe of firm-to-firm VAT invoices, (ii) the national business registry, and (iii) annual firm financial statements.

Firm-to-firm VAT invoices. The primary dataset is value-added tax (VAT) invoices collected by China's State Taxation Administration for 2014–2018. Each observation is a transaction between two registered entities and records unique tax identifiers for buyer and seller, the transaction amount, VAT amount, and the transaction date. Beginning in 2017, invoices additionally report product details (item names and codes), quantities, units/specifications, and unit prices. Between 2014 and 2018, the data cover approximately 16.1 billion transactions among around 16 million establishments. Because VAT invoices are required to claim input-tax credits and non-compliance is legally penalized, coverage of domestic inter-firm trade is very high.

The VAT database has three main limitations relevant to this study: (i) small firms below the small-scale threshold (average monthly sales under RMB 30,000, approximately USD 4,400) are exempt from issuing VAT invoices;[5] (ii) direct import/export transactions are not captured because one node of the transaction is not registered in China; and (iii) sales directly to final consumers do not generate VAT invoices. The first limitation, in particular, explains the discrepancy between the official shutdown counts (Table 1) and the number of firms in my analysis sample, which is restricted

---

[5] While these excluded entities represent a small fraction of aggregate economic output, they account for a large proportion of the total number of business entities in China. Common forms include neighborhood street vendors, small family-run convenience stores ("mom-and-pop" shops), and individual self-employed artisans, who typically operate in the informal or semi-formal economy and do not participate in the inter-firm production networks relevant to this study.

to firms large enough to be in the VAT system.

I use the VAT microdata for three purposes. First, I use the immediate pre-policy window (January–March 2017) to construct firm trade exposure matrices on both the purchase and sale sides. These are used to identify direct (first-order) and indirect (second-order) upstream and downstream links to Xiong'an polluters and to measure exposure to the shutdown campaign. Second, I aggregate transactions to the firm level to build the main outcomes—total purchases and total sales—used to track firm performance before and after the campaign. Third, I leverage the 2017–2018 product fields to construct measures of product and service substitutability, capturing the availability of alternative trading partners.

Business registry. To assign firm characteristics and enable linkage, I merge the VAT universe to the national business registration records maintained by the State Administration for Industry and Commerce (SAIC). At registration (and upon subsequent substantial changes), firms report address, shareholders, key management, registered capital, industry codes, and year of establishment. A firm identifier crosswalk links these records to the VAT data.

Firm annual financial statements. To track a longer post-campaign period beyond the VAT span (which ends in 2018), I extract firms' total revenues from annual self-reported financial statements. These statements, compiled by local registry offices and digitized after 2010, are available through 2020 and include assets, liabilities, revenues, profits, and taxes. Reporting is mandatory for active firms; failure to file for two consecutive years triggers license cancellation. I construct an annual panel from 2013 to 2020 using these data. Note that in addition to the sales to other firms covered by the VAT data, total revenue includes other sources of revenue, such as sales to final consumers, export revenue, and investment gains.

3 Sample, Measurement, and Specification

3.1 The Direct Effect of the Shutdown Campaign on Xiong'an Industrial Firms

Treatment Group A key identification challenge is the lack of official, firm-level data identifying which firms were subject to shutdown orders. Public announcements provided aggregate statistics on shutdowns rather than specific firm lists. Therefore, for my baseline analysis, I define the treatment group as all industrial firms located within the Xiong'an New Area prior to the policy announcement.

This comprehensive definition is justified for two primary reasons. First, the policy was implemented with unprecedented stringency as part of a top-down national strategy. The stated goal was to *eliminate* industrial pollution sources within the area, which was enforced through comprehensive and thorough inspections. Second, the effectiveness of this broad campaign is empirically validated by my results. As shown in Section 5, firms in this treatment group experienced an average reduction in both purchases and sales exceeding 50% following the policy's full implementation. This confirms that the campaign, on average, severely impacted these firms.

Control Group In principle, the control group could consist of all industrial firms outside Xiong'an with no trade links to the treated firms. However, this pool is impractically large (approximately 1,400 treated firms versus over 15 million potential controls), making estimation computationally burdensome. I therefore employ propensity score matching (PSM) to construct a comparable control sample. For each treated firm, I run a probit regression to select the 8 most similar untreated firms based on pre-campaign characteristics, including firm size, industry, and firm age.

I executed the matching algorithm allowing for replacement, implying that a control firm could theoretically be matched to multiple treated firms. However, given the vast pool of potential controls, overlaps were rare, affecting fewer than 100 firms. In instances where a control firm was selected for multiple treated firms, I exclusively assigned it to the treated firm with the relatively highest similarity score (compared to the match quality with other treated firms). This procedure ensures that each firm in the final control set C is unique and distinct from the treated set $\mathcal{S}$.

Specification Using this matched sample of treated firms ($\mathcal{S}$) and their counterparts (C), I estimate the direct effect of the pollution shutdown campaign using the following event-study specification:

$$y_{igsct} = \sum_{\tau \neq -1} \beta_\tau \cdot \text{XiongAn\_Ind}_i \times \text{Time}_\tau + \delta_i + \gamma_{sct} + \theta_{gt} + \varepsilon_{igsct} \tag{1}$$

where y_{igsct} is the log outcome (e.g., purchases or sales) for firm $i \in \mathcal{S} \cup C$ in matching group g, industry s, and province c, at time t. XiongAn_Ind_i is an indicator for treated industrial firms located in Xiong'an ($i \in \mathcal{S}$). Time_τ is an indicator for event time τ, relative to the policy launch (with $\tau = -1$ as the omitted base period).

The model includes a rich set of fixed effects. δ_i are firm fixed effects, which control for all time-invariant firm characteristics; estimates are thus identified from within-firm changes over time. γ_{sct} are industry-province-time fixed effects, which absorb any shocks common to firms in the same industry, province, and time period. θ_{gt} are matching-group-time fixed effects, which flexibly control for any unobserved trends specific to each matched group. ε_{igsct} is the error term. The coefficients of interest, β_τ, trace the dynamic treatment effect of the policy. Standard errors are clustered at the firm level to account for serial correlation.

3.2 The Propagation Effect on Upstream and Downstream Firms

Sample Construction for Propagation To measure the spillover effects, I exploit the variation in firms' pre-policy trade exposure to the shutdown firms. I construct my analysis samples by comparing the trade partners of the treated firms ($\mathcal{S}$) with the trade partners of the matched control firms ($\mathcal{C}$).

For the first-degree ($d = 1$) sample, the treatment group consists of firms with direct trade links—as suppliers or customers—to the Xiong'an industrial firms in $\mathcal{S}$, while the control group consists of firms with first-degree trade links to the matched controls in $\mathcal{C}$. The second-degree ($d = 2$) sample extends this logic one step further: the treatment group consists of suppliers-of-suppliers or customers-of-customers of $\mathcal{S}$, and the control group of the corresponding second-degree partners of $\mathcal{C}$. To isolate the pure second-degree effect, I apply an exclusivity rule: all firms with any first-degree exposure ($\text{UpExp}^{(1)} > 0$ or $\text{DownExp}^{(1)} > 0$) are excluded from the second-degree estimation sample.

A portion of firms, referred to as "inter-group" firms, had pre-policy trade links to both the treated Xiong'an firms ($\mathcal{S}$) and the matched control firms ($\mathcal{C}$). In my baseline, I assign these firms exclusively to the treatment group based on their exposure to Xiong'an. I test the robustness of this choice in Section 5.5.1.

Exposure Measurement I define a firm's trade exposure based on its pre-policy (pre-2017) VAT transactions with the Xiong'an industrial firms in $\mathcal{S}$. For a first-degree upstream supplier i, its exposure $\text{UpExp}_i^{(1)}$ is the share of total sales directed to firms in $\mathcal{S}$:

$$\text{UpExp}_i^{(1)} = \sum_{j \in \mathcal{S}} \frac{\text{sales}_{i \to j,\text{pre}}}{\sum_m \text{sales}_{i \to m,\text{pre}}}$$

For a first-degree downstream customer k, its exposure $\text{DownExp}_k^{(1)}$ is the share of total purchases sourced from firms in $\mathcal{S}$:[6]

$$\text{DownExp}_k^{(1)} = \sum_{j \in \mathcal{S}} \frac{\text{purchases}_{k \leftarrow j,\text{pre}}}{\sum_m \text{purchases}_{k \leftarrow m,\text{pre}}}$$

Second-degree exposure is defined recursively. A firm i's second-degree upstream exposure, $\text{UpExp}_i^{(2)}$, is the sales-weighted average of its direct buyers' first-degree upstream exposure:

$$\text{UpExp}_i^{(2)} = \sum_j r[i \to j] \cdot \text{UpExp}_j^{(1)}$$

---

[6] The denominator in both equations ($\sum_m$) represents the firm's total value of transactions with all trading partners in the economy during the pre-period, differing from the numerator which is restricted to the specific subset of treated firms.

where $r[i \to j]$ is the share of firm i's sales going to buyer j. The second-degree downstream exposure ($\text{DownExp}_i^{(2)}$) is calculated analogously using input shares and suppliers' first-degree downstream exposure.

Specification for Propagation I estimate the propagation effects for both first- and second-degree links using a unified event-study specification, where the effect is identified by the continuous exposure measure:

$$y_{isct} = \sum_{\tau \neq -1} \beta_\tau^{(d)} \cdot \text{Exp\_XA}_i^{(d)} \times \text{Time}_\tau + \delta_i + \gamma_{sct} + \mathbf{X}'_{i,\tau_0-1}\theta_t + \varepsilon_{isct} \tag{2}$$

where y_{isct} is the log outcome for firm i. The degree of propagation is denoted by $d \in \{1, 2\}$. The term $\text{Exp\_XA}_i^{(d)}$ is the pre-policy trade exposure of degree d (i.e., $\text{UpExp}_i^{(d)}$ or $\text{DownExp}_i^{(d)}$), normalized by its sample mean. The model includes firm fixed effects (δ_i) and industry-province-time fixed effects (γ_{sct}).

Critically, since I cannot use matching-group fixed effects in the propagation samples (as one firm may link to multiple treated firms), I instead control for a vector of firm characteristics from the pre-policy base period $\mathbf{X}_{i,\tau_0-1}$ (including firm size, age, and ownership type) interacted with time fixed effects (θ_t). This $\mathbf{X}'_{i,\tau_0-1}\theta_t$ term controls for any differential trends based on the same observable characteristics used in the initial PSM matching. The coefficients $\beta_\tau^{(d)}$ capture the average effect on firm i for a one-unit (mean-normalized) increase in its trade exposure of degree d.

3.3 Identification Assumptions

My research design employs a difference-in-differences (DID) framework to estimate both the direct impact of the regulation and its propagation through the supply chain. While both analyses rest on the fundamental assumption of parallel trends—that treated and control units would have evolved similarly in the absence of the shock—the specific identifying conditions and the role of my controls differ between the targeted firms and their trading partners.

Identification of the Direct Effect The estimation of the direct shutdown effect (Equation 1) relies on the assumption that the assignment of shutdown status is orthogonal to potential outcomes, conditional on my rich set of fixed effects. I argue this is plausible due to the specific institutional context and my rigorous control structure.

First, the inclusion of industry-province-time fixed effects (γ_{sct}) absorbs all time-varying shocks specific to a sector within a region. For example, if the entire steel industry in Hebei province faced a downturn due to global prices or local tax changes, this aggregate shock is controlled for, preventing it from confounding my estimates. Second, the matching-group-time fixed effects (θ_{gt}) play a distinct and critical role: they non-parametrically control for unobserved trends common to firms with similar pre-determined attributes. Since my matching was based on characteristics like size and age, these FEs ensure I am comparing treated firms to controls that are not just in the same industry, but are also on similar developmental trajectories. Ideally, if large, mature industrial firms nationwide faced a credit contraction during this period, this trend would be absorbed within the matched group, isolating the effect of the shutdown.

The primary threat to validity would be endogenous selection—for instance, if regulators targeted firms specifically because they were already financially distressed or shrinking. However, the Xiong'an campaign was a "top-down" political mandate driven by national strategic interests (the "Millennium Strategy" to relieve Beijing's non-capital functions) rather than local economic considerations. Unlike routine environmental enforcement, which is often triggered by local complaints or targets visibly struggling firms, the Xiong'an order was a blanket administrative clearance aimed at ecological restoration.

Identification of Propagation Effects The estimation of network spillovers (Equation 2) requires the assumption that a firm's pre-policy exposure to the shutdown targets is uncorrelated with unobserved trends in its own outcomes.

To satisfy this, I first control for industry-province-time fixed effects to sweep out any broad demand or supply shocks affecting specific sectors or regions. Furthermore, I include interactions of pre-policy firm characteristics (size, age, and ownership type) with time fixed effects. This is crucial because it separates the effect of network exposure from trends driven by firm attributes. For example, if suppliers of polluting firms tend to be smaller or state-owned, and such firms were on a slower growth path generally, my controls would account for these differential trends, leaving only the variation driven by the specific linkage to Xiong'an.

The exogeneity of this network variation is underpinned by the secrecy and suddenness of the policy. Because the Xiong'an plan was a closely guarded state secret prior to April 2017, supply chain partners could not have anticipated the shock. Consequently, the variation in exposure ($\text{UpExp}^{(d)}$ and $\text{DownExp}^{(d)}$) reflects historical trading patterns established long before the policy was conceived, rather than an endogenous response to regulatory risk. For instance, a risk-averse supplier might typically diversify away from a client facing imminent regulatory action; however, in this case, the risk of a total shutdown was effectively zero in the information set of market participants prior to the announcement. Thus, there was no anticipatory "severing of ties" that would otherwise bias my estimates.

Furthermore, for both the direct and propagation effects, the validity of my design is supported by the absence of pre-existing differential trends. If the parallel trends assumption were violated, I would expect to see divergence in outcomes prior to the policy announcement—for example, if Xiong'an firms were already on a different growth trajectory due to unobserved local factors, or if suppliers began severing ties in anticipation of regulatory risks. However, as I will demonstrate in

Section 5, the pre-policy coefficients in my event studies (for $\tau < -1$) are consistently small and statistically indistinguishable from zero. This lack of pre-trend provides strong visual and statistical evidence that the control group serves as a valid counterfactual for the treated firms.

3.4 Empirical Hypothesis

My empirical strategy aims to capture shocks propagating through the production network. To fix ideas, consider the stylized supply chain depicted in Figure 2. A steel mill in Xiong'an is shut down, severing its links to an upstream iron ore supplier and a downstream auto manufacturer.

Figure 2: Contextualizing Example: The Supply Chain

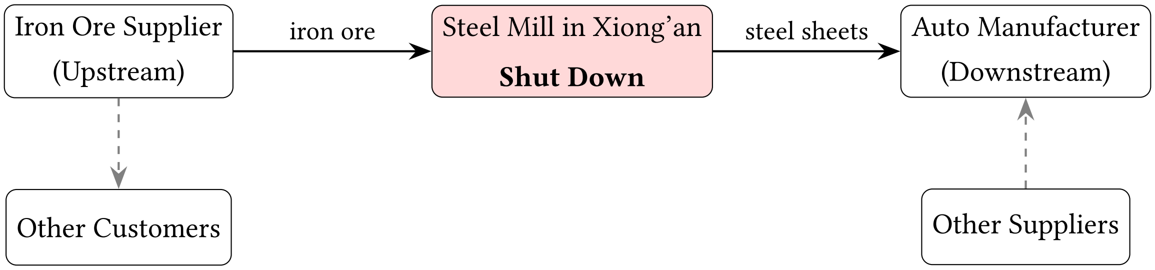

Notes: This figure illustrates the two primary channels of shock propagation. The upstream ore supplier loses a major customer (demand shock). The downstream auto manufacturer loses a key input source (supply shock). Dashed arrows represent remaining relationships with other partners. The key empirical question is whether the shock spills over to the ore supplier's other customers and the auto manufacturer's other suppliers—evidence of complementarity.

On the upstream side, the ore supplier loses a major customer, creating a demand shock. Its sales to the steel mill drop mechanically. If output complementarity is present ($\varphi > 0$), the loss of this customer erodes the ore supplier's productive capacity, reducing its ability to serve its remaining customers as well. On the downstream side, the auto manufacturer loses a key steel supplier, creating a supply shock. Its purchases from Xiong'an drop mechanically. If input complementarity is present ($\xi < 1$), the loss of one input forces proportional cutbacks across all inputs, reducing the auto manufacturer's purchases from its other suppliers too. I therefore expect firms with greater pre-policy trade exposure to Xiong'an ($\text{UpExp}^{(d)}$ or $\text{DownExp}^{(d)}$) to experience larger contractions in their own economic activity, implying negative coefficients $\beta_\tau^{(d)}$ for $\tau \geq 0$ in both the upstream and downstream samples.

These first-degree contractions in turn transmit new shocks outward. When the ore supplier scales back its own purchases, its suppliers (the second-degree upstream firms) face a demand shock. When the auto manufacturer reduces its own sales, its customers face a supply shock. The second-degree effects should be attenuated at the firm level ($|\beta^{(2)}| < |\beta^{(1)}|$), but the extensive margin grows rapidly: many more firms are affected at higher network degrees.

4 Descriptive Analysis

4.1 Xiong’an Industrial Firms

Figure 3 provides an overview of the Xiong’an industrial firms. Panels (a) and (b) show their industry composition by firm count and by sales share, respectively. “Rubber & Plastic Products” (453 firms), “Apparel” (220 firms), and “Metal Processing” (130 firms) are the most numerous. Metal Processing alone accounts for 38.6% of the group’s total pre-policy sales, followed by Rubber & Plastic Products (20.4%) and Apparel (10.3%). This makeup is consistent with the campaign’s environmental goals, as these industries are associated with heavy pollution and chemical emissions.

Panels (c) and (d) show network characteristics by firm size decile. A clear pattern emerges: larger firms have substantially more trade partners and conduct business over greater geographic distances. Firms in the smallest purchase decile average just over 1 supplier, compared to over 22 in the top decile. Average distance trends upward from 125km for the smallest to over 460km for the largest. This size–connectivity gradient proves critical for understanding the heterogeneous propagation documented below.

Figure 3: Xiong'an Polluting Firms: Industry Composition and Network Characteristics

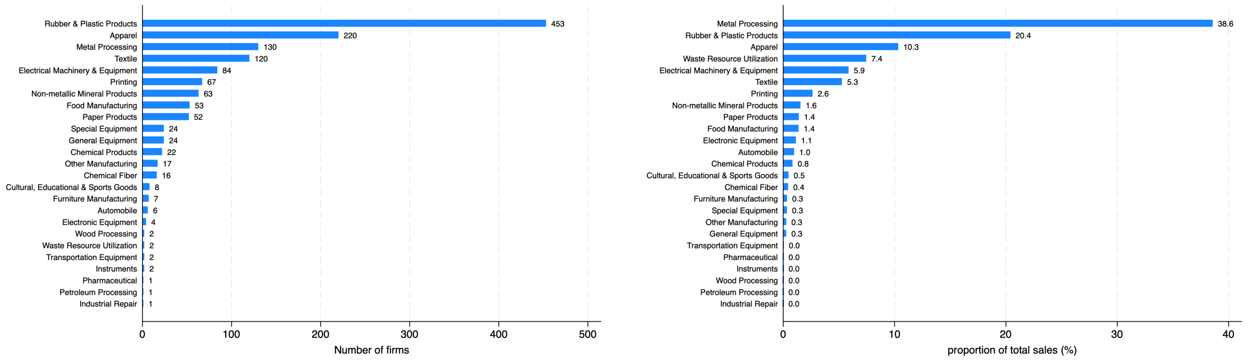

(a) Industry Composition (Number of Firms)

(b) Industry Composition (Share of Total Sales %)

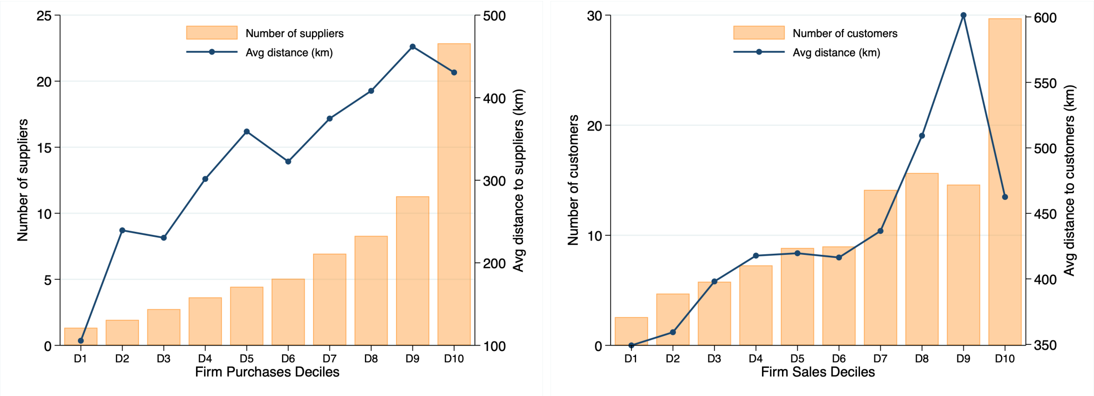

(c) Network by Firm Purchase Deciles

(d) Network by Firm Sales Deciles

Notes: Panels (a)–(b) show the industry composition of the 1,381 Xiong'an polluting firms. Panels (c)–(d) plot network characteristics by pre-policy size decile. Panel (c) sorts firms by total purchases and plots average number of suppliers (bars, left axis) and average distance to suppliers (line, right axis). Panel (d) sorts by total sales and plots average number of customers and distance to customers.

4.2 Network Characteristics

The industry composition of the affected production network is detailed in Appendix Figure A.1. As shown in Panels (a) and (b), "Wholesale & Retail Trade" firms are the most common first-degree partners (1,304 suppliers, 2,099 customers), highlighting the significant role of intermediaries in the supply chain. Beyond these intermediaries, the network is diverse. Upstream suppliers include significant clusters in "Textile" (431 firms), "Rubber & Plastic Products" (305 firms), and "Chemical Products" (295 firms), reflecting the inputs needed by the Xiong'an industrial firms. Downstream customers are similarly varied, with large numbers in "Apparel" (617 firms), "Research & Development" (585 firms), and "Electrical Machinery & Equipment" (459 firms), showing that the outputs served a wide range of subsequent manufacturing and service activities. As shown in Panels (c) and (d), this pattern continues at the second degree, where intermediaries remain dominant but are followed by a wide array of other manufacturing and service sectors.

I also examine the spatial distribution of firms in my samples, shown in Appendix Figure A.2. Panel (a) shows the Xiong'an polluting firms are, by definition, highly concentrated in the Xiong'an area. Panel (b) shows their first-degree trading partners; while a large proportion is located in the surrounding Jing-Jin-Ji area, many firms are also located farther away across the country. As I move to the second degree in Panel (c), the network becomes even more spatially diffuse. Panels (d) through (f) demonstrate that my matched control firms and their trading partners are geographically dispersed and largely distant from the Xiong'an area, reflecting the national scope of the potential control pool.

Appendix Figure A.3 visualizes the distributions of the exposure measures. The first-degree distributions are highly right-skewed: the vast majority of firms have a very small exposure (less than 10%), but a long tail of firms are heavily dependent on Xiong'an industrial firms, with exposures approaching 100%. This heterogeneity motivates my use of a continuous exposure variable rather than a simple binary indicator. The second-degree exposure distributions confirm rapid attenuation; almost all firms have an exposure level that is effectively zero, supporting the small mean values in Table B.1 and my decision to not pursue higher-order effects.

4.3 Summary Statistics of Working Samples

Table 2 presents summary statistics for the pre-policy period for firms in the direct effect and first-degree propagation samples. Panel A compares the 1,381 Xiong'an polluting firms ($\mathcal{S}$) with their 10,516 matched control firms ($\mathcal{C}$). The propensity score matching appears effective: the two groups are highly comparable in terms of average log purchases (13.95 vs. 14.09) and log sales (14.08 vs. 14.29), supporting the validity of my control group for estimating the direct effect.

Panels B and C describe the first-degree propagation samples. Panel B shows that 7,748 upstream suppliers to Xiong'an polluting firms had an average of 5% of their total sales directed to the soon-to-be-shutdown Xiong'an firms ($\text{UpExp}^{(1)}$). These suppliers are substantially larger than control suppliers, with average purchases of 184.7 million RMB compared to 63.3 million RMB for control suppliers. They also have extensive customer networks, averaging 329 customers. Panel C shows that 14,103 downstream customers had an average of 7% of their total purchases from Xiong'an firms ($\text{DownExp}^{(1)}$). These customers average 120 suppliers.

Table B.1 presents summary statistics for the second-degree propagation samples. As I move further from the initial shock, two patterns emerge. First, sample sizes expand dramatically: 331,755 second-degree upstream firms and 960,640 second-degree downstream firms. Second, the average exposure levels are dramatically smaller: the mean second-degree upstream and downstream exposures are only 0.07% and 0.06%, respectively. This rapid attenuation of exposure motivates my decision to focus on first- and second-degree effects without pursuing higher-order propagation.

Table 2: Summary Statistics of Pre-Policy Firm Characteristics: Direct & 1st Degree Propagation Effects

	(1)					(2)				
Variables	Mean	SD	P5	P95	N	Mean	SD	P5	P95	N
Panel A: Sample of the direct effect	Industrial Firms in XiangAn					Control Firms				
Purchase (Million RMB)	10.30	68.89	0.01	32.69	1,381	10.58	63.42	0.03	38.08	10,516
Sale (Million RMB)	10.71	72.05	0.03	33.76	1,381	11.38	62.94	0.07	40.12	10,516
Log Purchase	13.95	3.19	9.47	17.30	1,381	14.09	2.71	10.43	17.46	10,516
Log Sale	14.08	3.38	10.35	17.33	1,381	14.29	2.88	11.10	17.51	10,516
Age	10.35	5.47	3.00	21.00	1,381	10.35	5.48	3.00	20.00	10,516
SOE	0.00	0.05	0.00	0.00	1,381	0.01	0.09	0.00	0.00	10,516
Distance to the Xiong'an Centrioid (km)	15.94	7.44	5.29	31.11	1,381	1056.67	529.02	133.08	1823.56	10,516
Panel B: Sample of the 1st degree upstream propagation	1st Degree Suppliers of Xiong'an Industrial Firms					1st Degree Suppliers of the Control Firms				
Purchase (Million RMB)	184.70	1284.52	0.30	576.39	7,748	63.33	832.48	0.20	178.43	144,973
Sale (Million RMB)	205.76	1387.62	0.56	631.13	7,748	70.91	965.31	0.39	189.83	144,973
Log Purchase	16.30	2.49	12.60	20.17	7,748	15.42	2.28	12.20	19.00	144,973
Log Sale	16.60	2.15	13.24	20.26	7,748	15.72	1.90	12.88	19.06	144,973
Age	11.24	6.19	3.00	23.00	7,733	10.82	6.00	3.00	22.00	144,735
SOE	0.04	0.21	0.00	0.00	7,748	0.03	0.17	0.00	0.00	144,973
Distance to the Xiong'an Centrioid (km)	625.71	558.98	44.45	1777.10	7,748	1059.13	504.09	132.72	1821.01	144,973
# of Customers	328.59	1379.07	7	1216	7,748	153.52	516.11	6	568	144,973
# of Other Customers	326.91	1378.85	6	1214	7,748	152.10	515.25	5	566	144,973
% Sales to Xiong'an Industrial Firms	0.05	0.13	0.00	0.23	7,748	-	-	-	-	144,973
Panel C: Sample of the 1st degree downstream propagation	1st Degree Customers of Xiong'an Industrial Firms					1st Degree Customers of the Control Firms				
Purchase (Million RMB)	128.48	2033.88	0.38	364.94	14,103	80.89	708.62	0.35	244.69	165,383
Sale (Million RMB)	126.63	2148.60	0.39	322.42	14,103	86.99	858.25	0.37	245.21	165,383
Log Purchase	16.12	2.11	12.86	19.72	14,103	15.85	2.01	12.75	19.32	165,383
Log Sale	16.03	2.29	12.88	19.59	14,103	15.83	2.26	12.82	19.32	165,383
Age	11.33	6.44	3.00	23.00	14,080	10.99	6.05	3.00	22.00	165,153
SOE	0.06	0.23	0.00	1.00	14,103	0.03	0.18	0.00	0.00	165,383
Distance to the Xiong'an Centrioid (km)	564.04	538.17	51.94	1769.06	14,103	1064.98	501.84	146.58	1824.78	165,383
# of Other Suppliers	119.77	231.46	6	437	14,103	105.07	179.34	7	366	165,383
# of Non-Target Suppliers	118.57	231.38	5	436	14,103	103.87	179.16	6	364	165,383
% Purchases from Xiong'an Industrial Firms	0.07	0.16	0.00	0.40	14,103	-	-	-	-	165,383

Notes: This table presents summary statistics for the pre-policy period (October 2016–March 2017). Panel A shows the 1,381 Xiong'an polluting firms ($\mathcal{S}$) and their 10,516 matched control firms ($\mathcal{C}$). Panels B and C show the first-degree upstream (suppliers) and downstream (customers) samples, respectively. "Log Purchase" and "Log Sale" are measured in the pre-policy period. "% Sales to Xiong'an Industrial Firms" (Panel B) and "% Purchases from Xiong'an Industrial Firms" (Panel C) are the first-degree exposure measures. Distance is measured from each firm to the Xiong'an centroid in kilometers.

Figure 4: Trends of Xiong'an Industrial Firms vs. Control Firms over Time

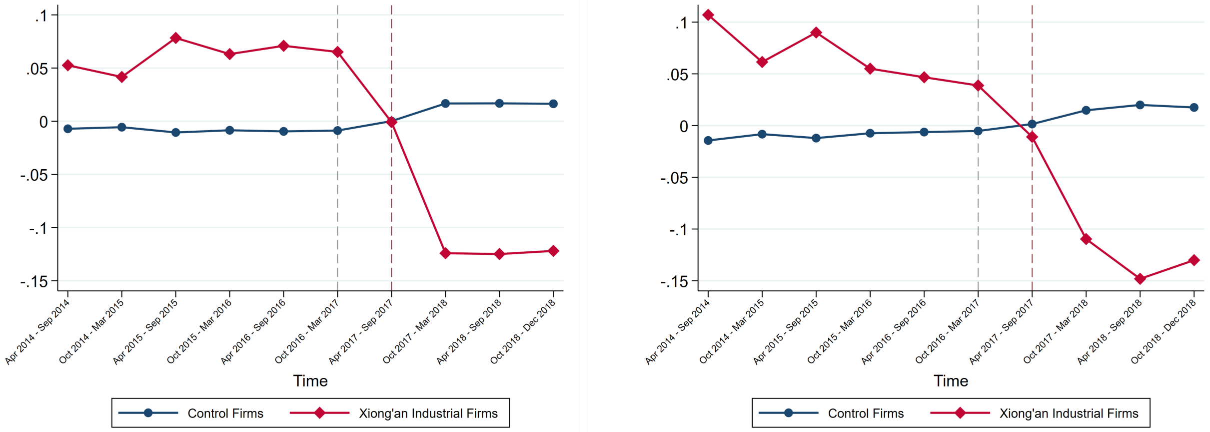

(a) Mean of Residualized Log Purchase

(b) Mean of Residualized Log Sale

Notes: This figure plots the mean of residualized log purchases (Panel a) and log sales (Panel b) for Xiong'an industrial firms (treated) and their matched control firms over time. Residuals are obtained from a regression of log outcomes on firm fixed effects and industry-province-time fixed effects. The vertical dashed lines indicate the "Partial Implementing" (April 2017) and "Full Implementing" (October 2017) phases. If control firms were benefiting from the exit of Xiong'an firms (e.g., capturing their market share), I would expect the control group's residualized outcomes to increase after the policy. Instead, the control group remains flat while only the treated group declines, consistent with the shutdown shock operating through supply chain linkages rather than product market competition.

5 Empirical Results

In this section, I present the empirical results, beginning with the direct impact on the targeted Xiong'an polluting firms and then tracing the propagation of this shock through the first- and second-degree production network.

5.1 Direct Effect

I first validate that the shutdown campaign had a significant and immediate impact on the targeted firms. Figure 5 plots the event-study coefficients (β_τ from Equation 1) for the log purchases and log sales of Xiong'an industrial firms relative to their matched controls. The coefficients for all pre-policy periods are small and statistically indistinguishable from zero, confirming the parallel trends assumption.

Following the policy's announcement in April 2017 ("Partial Implementing"), a negative effect begins to emerge. This effect becomes dramatically larger and statistically significant after October 2017 ("Full Implementing"), when the shutdowns were broadly executed. The magnitude of the collapse is stark: firms' purchases and sales both plummet, remaining at this new, lower level for the remainder of my sample period.

Table 3 corroborates these findings using the static DID specification. My baseline specification in column (2) shows that during the "Full Implementing" phase, Xiong'an firms experienced a massive 80.2 log-point drop in purchases, which translates to a 55.2% decline (i.e., $100 \times (e^{-0.802} - 1)$). Similarly, column (4) shows a 75.5 log-point drop in sales, implying a 53.0% decline. These results confirm that the shutdown campaign was an immediate, severe, and persistent negative shock to the targeted firms, validating my treatment definition and providing a strong basis for examining propagation effects.

Table 3: Static DID Results: Direct Effect (VAT Data)

	log_purchase		log_sale	
	(1)	(2)	(3)	(4)
XiongAn × Partial	-0.2625	-0.2555	-0.3234	-0.2591
	(0.2110)	(0.2139)	(0.2236)	(0.2260)
XiongAn × Full	-0.7757***	-0.8016***	-0.8181***	-0.7549***
	(0.2495)	(0.2591)	(0.2776)	(0.2884)
Observations	107,870	106,120	107,870	106,120
R^2	0.6705	0.7289	0.5827	0.6598
Firm FE	X	X	X	X
Time-Province-Industry FE	X	X	X	X
Time-Group FE		X		X

Notes: This table presents the static DID results for the direct effect on Xiong'an polluting firms, corresponding to Equation 1. The dependent variables are log purchases (Columns 1–2) and log sales (Columns 3–4). The sample includes Xiong'an polluting firms and their matched controls. "Partial" refers to the partial implementation period (April–September 2017) and "Full" refers to the full implementation period (October 2017–December 2018). Columns (2) and (4) include Time-Group fixed effects and are my baseline specifications. Standard errors clustered at the firm level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Figure 5: Event Study Results: Direct Effect on Xiong'an Polluting Firms

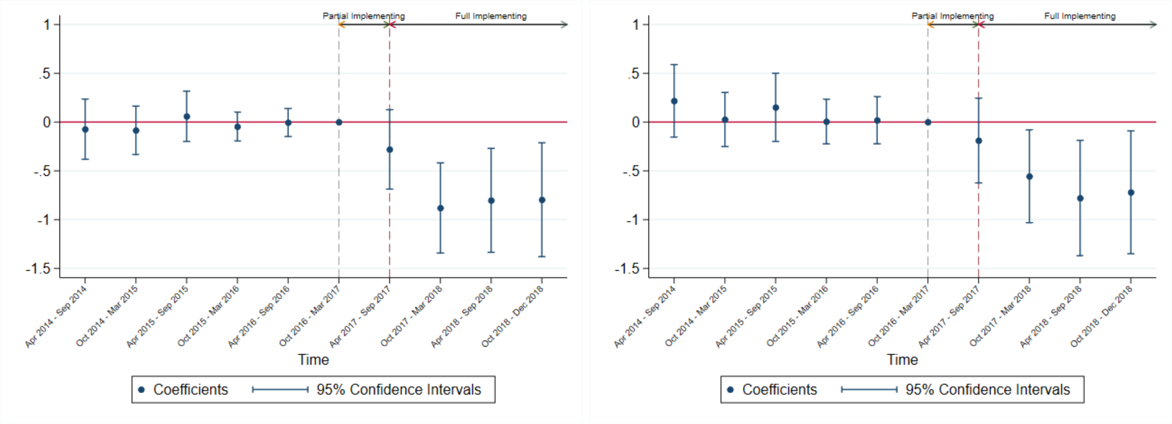

(a) Log Purchases (b) Log Sales

Notes: This figure plots the event-study coefficients β_τ from Equation 1, using semi-annual data. The omitted reference period is October 2016–March 2017 ($\tau = -1$). The specification includes firm fixed effects, industry-province-time fixed effects, and matching-group-time fixed effects. The vertical lines plot the 95% confidence intervals. Standard errors are clustered at the firm level. The vertical dashed lines indicate the "Partial Implementing" (April 2017) and "Full Implementing" (October 2017) phases.

5.2 First-Degree Propagation Effects

Having established the severe direct shock, I now test whether this shock propagated to the direct suppliers and customers of the shutdown firms. The event-study coefficients for first-degree upstream suppliers are plotted in Figure 6 (top row). The results show clear evidence of a negative demand shock. As with the direct effect, the pre-policy coefficients are flat and near zero. Following the full implementation phase, suppliers with higher pre-policy sales exposure to Xiong'an firms experience a statistically significant decline in their own log purchases (Panel a) and log sales (Panel b).

The static DID results in Table 4 (Columns 1–2) quantify this impact. Since the exposure variable is mean-normalized, the coefficients represent the effect for a supplier with the sample-mean exposure (5% pre-policy sales share to Xiong'an). The estimates show a 7.7% drop in purchases and an 11.8% drop in sales for mean-exposed suppliers. This suggests that the collapse in demand from their Xiong'an buyers forced these suppliers to contract their own operations, reducing their purchases of new inputs.

Figure 6 (bottom row) shows the corresponding results for first-degree downstream customers. I again find strong evidence of propagation. After the full implementation phase, firms that relied more heavily on Xiong'an firms for their inputs saw a sharp and persistent decline in their own activity. Specifically, Table 4 (Columns 3–4) shows that a downstream customer with the sample-mean exposure (7% pre-policy purchase share from Xiong'an) experienced a 9.5% drop in their own purchases and a 7.9% contraction in sales during the full implementation period. This supports my hypothesis that these firms were unable to perfectly substitute their lost suppliers, leading to a contraction in their own output.

These first-degree results reveal the initial transmission of the shock. For upstream suppliers, the reduction in their own sales is the direct demand shock. The accompanying drop in their purchases, however, indicates that they scaled back their own operations, thereby transmitting a new, negative demand shock to their suppliers (the second-degree upstream firms). Similarly, for downstream customers, the drop in their purchases (the direct supply shock) led to a contraction in their sales, transmitting a negative supply shock to their own customers (the second-degree downstream firms).

Figure 6: Event Study Results: First-Degree Propagation Effects

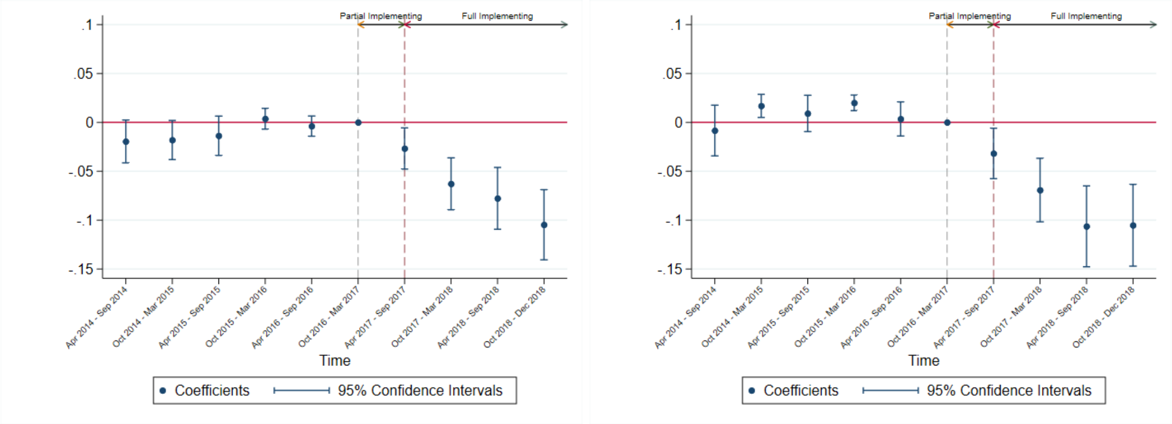

(a) Upstream (Log Purchases)

(b) Upstream (Log Sales)

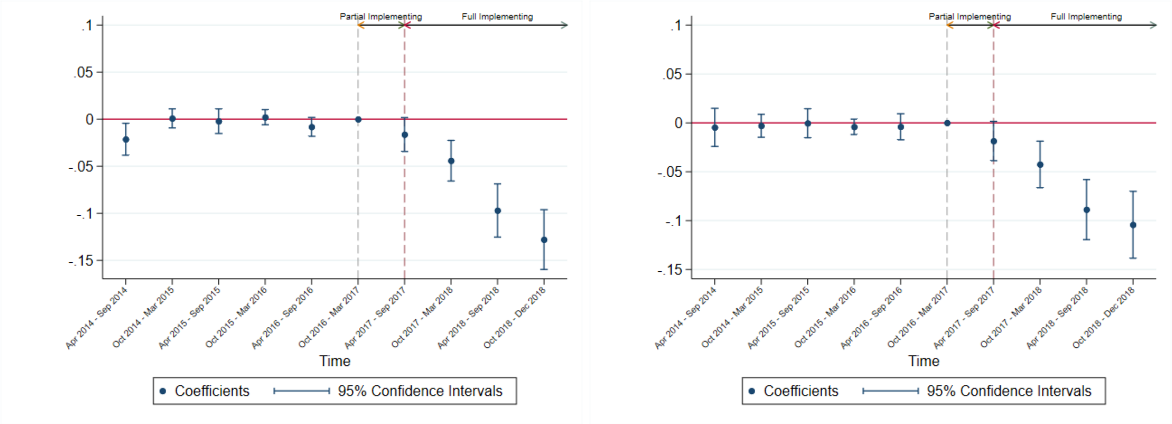

(c) Downstream (Log Purchases)

(d) Downstream (Log Sales)

Notes: This figure plots the event-study coefficients $\beta_\tau^{(1)}$ from Equation 2, using semi-annual VAT data. Panels (a) and (b) use the 1st-degree upstream sample; Panels (c) and (d) use the 1st-degree downstream sample. The omitted reference period is October 2016–March 2017 ($\tau = -1$). The specification includes firm fixed effects, industry-province-time fixed effects, and pre-policy firm characteristics interacted with time fixed effects. The vertical lines plot the 95% confidence intervals. Standard errors are clustered at the firm level. The vertical dashed lines indicate the “Partial Implementing” (April 2017) and “Full Implementing” (October 2017) phases.

Table 4: Static DID Results: First-Degree Propagation Effects

	Sample: 1st degree upstream firms		Sample: 1st degree downstream firms	
	log_purchase	log_sale	log_purchase	log_sale
	(1)	(2)	(3)	(4)
$\text{UpExp}^{(1)} \times \text{Partial}$	-0.0078	-0.0438^{***}		
	(0.0115)	(0.0132)		
$\text{UpExp}^{(1)} \times \text{Full}$	-0.0771^{***}	-0.1180^{***}		
	(0.0153)	(0.0197)		
$\text{DownExp}^{(1)} \times \text{Partial}$			-0.0121	-0.0172
			(0.0090)	(0.0105)
$\text{DownExp}^{(1)} \times \text{Full}$			-0.0947^{***}	-0.0794^{***}
			(0.0133)	(0.0147)
Observations	1,509,850	1,509,850	1,777,280	1,777,280
R^2	0.7973	0.7266	0.8062	0.6622
Firm FE	X	X	X	X
Time-Province-Industry FE	X	X	X	X
Pre-Policy Controls × Time FE	X	X	X	X
Mean 1st degree % sales to XiongAn	0.05	0.05		
Mean 1st degree % purchases from XiongAn			0.07	0.07

Notes: This table presents the static DID results for the first-degree propagation effects, corresponding to Equation 2 (for $d = 1$). Columns (1)–(2) use the 1st-degree upstream sample; Columns (3)–(4) use the 1st-degree downstream sample. $\text{UpExp}^{(1)}$ is the mean-normalized pre-policy sales share to Xiong'an firms. $\text{DownExp}^{(1)}$ is the mean-normalized pre-policy purchase share from Xiong'an firms. "Partial" refers to the partial implementation period (April–September 2017) and "Full" refers to the full implementation period (October 2017–December 2018). All specifications include firm fixed effects, industry-province-time fixed effects, and pre-policy firm characteristics (size, age, ownership) interacted with time fixed effects. Standard errors clustered at the firm level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

5.3 Second-Degree Propagation Effects

I next test whether the shock propagated even further, to the suppliers of suppliers and customers of customers. The event-study plots in Figure 7 show that the propagation continues, but with significant attenuation. For second-degree upstream suppliers, I see small, statistically significant drops in both purchases (Panel a) and sales (Panel b) beginning in the full implementation phase.

The static DID results in Table 5 quantify this attenuation: the estimates show approximately a 0.4% drop in purchases and 0.7% drop in sales for mean-exposed second-degree upstream firms. For second-degree downstream customers, the effects are even smaller: a 0.07% drop in purchases (significant at 5%) and a statistically insignificant 0.04% drop in sales.

While these effects are statistically significant, their economic magnitude on the intensive margin attenuates sharply—the second-degree effects are roughly 10–15 times smaller than the first-degree effects. This attenuation at the firm level, however, is counter-balanced by a rapid expansion of the extensive margin. As shown in Table B.1, the number of affected second-degree suppliers

(331,755) and customers (960,640) is substantially larger than the number of first-degree firms (7,748 and 14,103). Consequently, even with small individual impacts, the large number of affected firms implies that the aggregate economic loss remains significant at this stage of the network.

The contraction in purchases among second-degree upstream firms and sales among second-degree downstream firms indicates the potential for even further propagation to third-degree (and higher-order) firms. However, given the rapid attenuation of the per-firm effect size and the sample restrictions that arise from the exponential growth of trade partners, I do not formally identify these higher-order effects.

Figure 7: Event Study Results: Second-Degree Propagation Effects

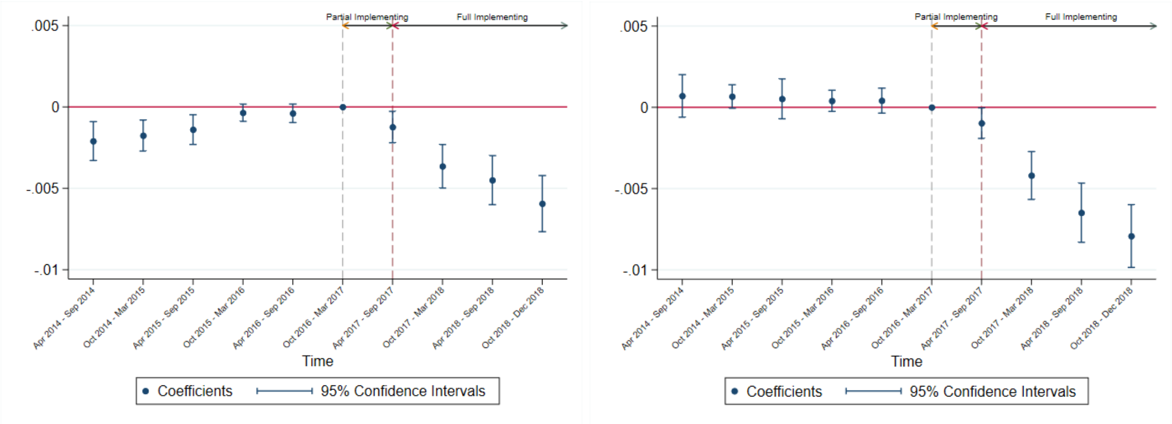

(a) 2nd-Degree Upstream (Log Purchases) (b) 2nd-Degree Upstream (Log Sales)

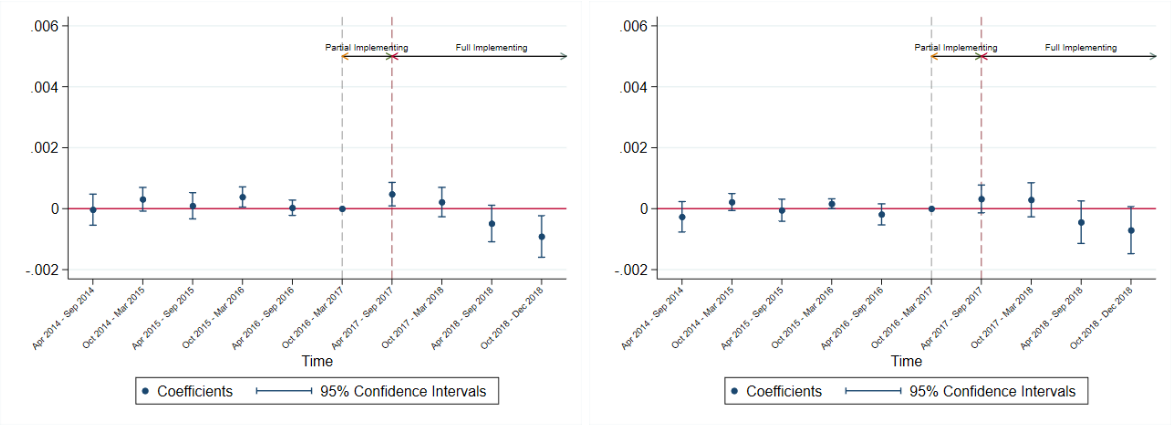

(c) 2nd-Degree Downstream (Log Purchases) (d) 2nd-Degree Downstream (Log Sales)

Notes: This figure plots the event-study coefficients $\beta_\tau^{(2)}$ from Equation 2 for the 2nd-degree samples, using semi-annual VAT data. Panels (a) and (b) show results for the upstream sample; Panels (c) and (d) show results for the downstream sample. The omitted reference period is October 2016–March 2017 ($\tau = -1$). The specification includes firm fixed effects, industry-province-time fixed effects, and pre-policy firm characteristics interacted with time fixed effects. The vertical lines plot the 95% confidence intervals. Standard errors are clustered at the firm level. The vertical dashed lines indicate the "Partial Implementing" (April 2017) and "Full Implementing" (October 2017) phases.

Table 5: Static DID Results: Second-Degree Propagation Effects

	Sample: 2nd degree upstream firms		Sample: 2nd degree downstream firms	
	log_purchase	log_sale	log_purchase	log_sale
	(1)	(2)	(3)	(4)
$\text{UpExp}^{(2)} \times$ Partial	-0.0002	-0.0014^{***}		
	(0.0005)	(0.0005)		
$\text{UpExp}^{(2)} \times$ Full	-0.0040^{***}	-0.0072^{***}		
	(0.0007)	(0.0009)		
$\text{DownExp}^{(2)} \times$ Partial			0.0003	0.0003
			(0.0002)	(0.0002)
$\text{DownExp}^{(2)} \times$ Full			-0.0007^{**}	-0.0004
			(0.0003)	(0.0003)
Observations	15,094,230	15,094,230	20,621,670	20,621,670
R^2	0.7561	0.6477	0.7376	0.6128
Firm FE	X	X	X	X
Time-Province-Industry FE	X	X	X	X
Pre-Policy Controls × Time FE	X	X	X	X
Mean 2nd degree % sales to XiongAn	0.0007	0.0007		
Mean 2nd degree % purchases from XiongAn			0.0006	0.0006

Notes: This table presents the static DID results for the second-degree propagation effects, corresponding to Equation 2 (for $d = 2$). Columns (1)–(2) use the 2nd-degree upstream sample; Columns (3)–(4) use the 2nd-degree downstream sample. $\text{UpExp}^{(2)}$ is the mean-normalized pre-policy second-degree sales share to Xiong'an firms. $\text{DownExp}^{(2)}$ is the mean-normalized pre-policy second-degree purchase share from Xiong'an firms. "Partial" refers to the partial implementation period (April–September 2017) and "Full" refers to the full implementation period (October 2017–December 2018). All specifications include firm fixed effects, industry-province-time fixed effects, and pre-policy firm characteristics (size, age, ownership) interacted with time fixed effects. Standard errors clustered at the firm level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

5.4 Persistence of Effects: Annual Revenue Data

While the semi-annual VAT data allows me to trace the precise timing and transmission of the shock, I turn to annual firm financial statement data, which extends to 2020, to examine longer-run persistence. Figure 8 plots the event-study coefficients. The direct effect on Xiong'an firms was severe and lasting: a 52.5% decline in revenue with no recovery through 2020. The first-degree propagation effects are similarly persistent: upstream suppliers experienced a 5.0% revenue decline and downstream customers a 4.3% decline, both significant for at least three years. Small but persistent negative effects appear for the second-degree samples as well.

Figure 8: Event Study Results: Effects on Annual Log Revenue (2013–2020)

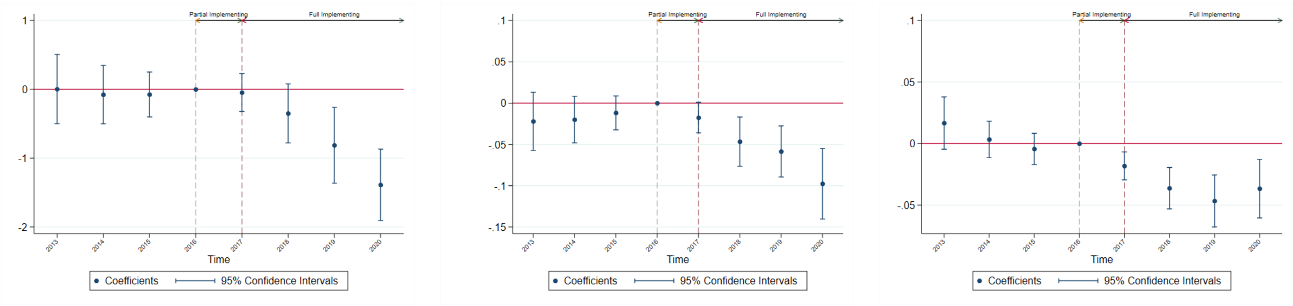

(a) Direct Effect

(b) 1st-Degree Upstream

(c) 1st-Degree Downstream

(d) 2nd-Degree Upstream

(e) 2nd-Degree Downstream

Notes: This figure plots event-study coefficients on annual log revenue. The omitted reference period is 2016 ($\tau = -1$). Panel (a) shows the direct effect (Eq. 1). Panels (b)–(e) show the propagation effects (Eq. 2). Standard errors are clustered at the firm level.

The persistence of these spillovers underscores network stickiness: firms face substantial frictions in forming new trade relationships to replace lost partners. Appendix Table B.2 reports the corresponding static DID estimates, confirming these patterns. I interpret the 2020 estimates with caution, as the onset of COVID-19 introduced a concurrent shock and higher rates of missing data.

5.5 Robustness Checks

To ensure my results are not driven by sample selection or confounding factors, I perform two sets of robustness checks.

5.5.1 Alternative Sample: "Inter-Group" Firms

A potential concern in my propagation analysis is the treatment of "inter-group" firms—those that had pre-policy trade links to both the treated Xiong'an firms ($\mathcal{S}$) and the matched control firms ($\mathcal{C}$). In my baseline specification (Section 3), I exclusively assign these firms to the treatment group based on their exposure to Xiong'an firms. In Appendix Table B.3, I test two alternative sample definitions: excluding these firms entirely, and duplicating them by assigning them to both groups. The results remain quantitatively similar and statistically significant across both specifications.

5.5.2 Potential Confounding Spillovers

The estimated coefficients in my main specifications capture the net effect of the shutdown campaign on exposed firms. This net effect includes the primary negative demand and supply shocks, but it could also theoretically include positive spillovers from within-firm reallocation or reduced product-market competition. I argue that my estimates are dominated by the primary negative shock, as these other potential channels appear to be trivial in my setting.

For within-firm reallocation, a multi-plant parent company could respond to the shutdown of its Xiong'an plant by shifting production to a plant outside the area, creating a positive spillover within the corporate group. I use linked shareholder information to explore this channel and find that, among the 1,381 Xiong'an polluting plants in my sample, only 9 are part of a larger multi-plant firm. Given this, any within-firm reallocation effect is trivial in my setting.

For product-market competition, my matching strategy selects control firms from the same industry as the treated firms, so the shutdown could create a positive demand shock for control-group competitors. I argue this effect is limited, as Xiong'an firms were generally small and the national control pool is geographically dispersed. Figure 4 provides further evidence: if control firms were benefiting from reduced competition, their residualized outcomes should increase after the policy. Instead, the control group's outcomes remain flat, with the divergence driven entirely by the treated group's decline. Alternative restrictions on the control pool—excluding geographically proximate firms or firms sharing pre-policy trade partners with their matched treated firm—yield similar results.

6 Heterogeneity

Having established the average propagation effects, I now explore which firm characteristics buffer or amplify these spillovers. I focus on the first-degree upstream and downstream samples and augment my baseline specification (Equation 2) by interacting the exposure measure with a time-invariant, pre-policy firm characteristic, Z_i. This allows me to move from an average effect to a differential effect based on observable firm types.

I estimate the following static DID specification, which includes interactions for both the partial and full implementation phases:

$$\begin{aligned} y_{isct} = & \sum_{p \in \{\text{Part}, \text{Full}\}} \left(\alpha_{1,p} \cdot \text{Exp\_XA}_i^{(1)} \times \text{Time}_p + \alpha_{2,p} \cdot \text{Exp\_XA}_i^{(1)} \times \text{Time}_p \times Z_i \right) \\ & + \mathbf{X}'_{i,\tau_0 - 1} \theta_t + \delta_i + \gamma_{sct} + \varepsilon_{isct} \end{aligned} \tag{3}$$

where Time_p is an indicator for either the partial implementation period (p = Part) or the full implementation period (p = Full). $\text{Exp\_XA}_i^{(1)}$ is the mean-normalized first-degree upstream ($\text{UpExp}^{(1)}$) or downstream ($\text{DownExp}^{(1)}$) exposure. As in my baseline model, I include firm fixed effects (δ_i), industry-province-time fixed effects (γ_{sct}), and pre-policy firm controls vector $\mathbf{X}_{i,\tau_0-1}$ interacted with time fixed effects (θ_t).

The coefficients $\alpha_{1,p}$ capture the baseline propagation effect for firms with $Z_i = 0$ during period p, while $\alpha_{2,p}$ captures the differential effect for firms with that characteristic. For brevity, in the following tables I report only the coefficients for the full implementation period ($\alpha_{1,\text{Full}}$ and $\alpha_{2,\text{Full}}$), as this is when the shock fully materialized.

My central finding is that firm size emerges as the dominant driver of heterogeneous propagation effects. Small firms experience substantially larger losses relative to their baseline, while large firms are better able to absorb supply chain shocks. In Section 7, I explore the mechanisms underlying this pattern.

6.1 Heterogeneity by Firm Size

I proxy for firm size using pre-policy total sales (inputs) and create a binary indicator Z_i that equals one if the firm is "Large" (above the sample median) and zero otherwise. My prior is that larger firms possess greater financial and operational resources, which should allow them to absorb shocks and find alternative partners more easily.

Table 6 presents the results. For the baseline group of smaller firms (below median size), the shock is severe: upstream sales drop by 13.0% and downstream purchases drop by 10.7% for mean-exposed small firms. However, large firms experience a significantly more positive outcome. In the upstream sample, the interaction term of 13.1% implies that large firms saw a differential effect that almost fully offsets the baseline decline—the implied net effect on large firms is only 0.1% (essentially zero). In the downstream sample, the interaction term of 13.2% for purchases similarly neutralizes the shock: the implied net effect on large firms is actually positive 2.5%, meaning large downstream firms show no significant decline. These results indicate that scale is a critical factor in a firm's ability to survive supply chain disruptions.

Table 6: Heterogeneous Propagation Effects by Firm Size

	Sample: 1st degree upstream firms		Sample: 1st degree downstream firms	
	log_purchase	log_sale	log_purchase	log_sale
	(1)	(2)	(3)	(4)
$\text{UpExp}^{(1)} \times \text{Full}$	-0.0861***	-0.1299***		
	(0.0163)	(0.0212)		
$\text{UpExp}^{(1)} \times \text{Full} \times \text{Large}$	0.0988**	0.1310***		
	(0.0405)	(0.0461)		
$\text{DownExp}^{(1)} \times \text{Full}$			-0.1070***	-0.0926***
			(0.0143)	(0.0157)
$\text{DownExp}^{(1)} \times \text{Full} \times \text{Large}$			0.1321***	0.1410***
			(0.0274)	(0.0334)
Observations	1,509,850	1,509,850	1,777,280	1,777,280
R^2	0.7973	0.7266	0.8063	0.6622
Firm FE	X	X	X	X
Time-Province-Industry FE	X	X	X	X
Pre-Policy Controls × Time FE	X	X	X	X
Implied effect on large firms	0.0127	0.0011	0.0251	0.0484
	(0.0371)	(0.0409)	(0.0240)	(0.0302)
Mean 1st degree % sales to XiongAn	0.05	0.05		
Mean 1st degree % purchases from XiongAn			0.07	0.07

Notes: This table presents heterogeneous effects of the first-degree propagation by firm size, corresponding to Equation 3. "Large" is an indicator for firms with above-median pre-policy total sales. The baseline coefficient captures the effect for small firms (below median), while the interaction term captures the differential effect for large firms. A positive interaction coefficient indicates that large firms experience smaller losses (i.e., are better buffered against the shock). The "Implied effect on large firms" row reports the sum of the baseline and interaction coefficients. All specifications include firm fixed effects, industry-province-time fixed effects, and pre-policy firm characteristics interacted with time fixed effects. Standard errors clustered at the firm level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

6.2 Heterogeneity by Firm Sector

I also explore heterogeneous propagation effects across sectors. The results, reported in Appendix Table B.4, show no statistically significant differences across industrial, trade intermediary, and service sectors. These sector-based results should be interpreted with caution, however, as sector is partially confounded with firm size: industrial firms in my sample tend to be larger on average than service firms. Given the dominant role of firm size documented above, the apparent absence of sectoral differences likely reflects the fact that firm size subsumes any sector-specific channel.

6.3 Aggregate Network Propagation Losses: A Back-of-Envelope Calculation

The size heterogeneity documented above has a crucial implication for quantifying the aggregate economic cost of the shutdown campaign. Since small firms experience large percentage losses but account for only about 1% of total group sales, while large firms—holding 99% of sales—show statistically zero propagation effects, aggregating the pooled regression coefficient across all firms would vastly overstate the true spillover. I therefore compute aggregate losses that account for size-varying elasticities, size-varying trade exposure, and the extreme skewness of the sales distribution. Since the implied propagation effects on large firms are statistically indistinguishable from zero in all specifications (Table 6), I restrict the aggregation to small firms.[7] The detailed methodology and step-by-step calculation are reported in Appendix C.

The aggregate first-degree spillover loss, focusing exclusively on small firms, amounts to approximately 46% of the direct loss for upstream suppliers and 39% of the direct loss for downstream customers, for a combined spillover-to-direct ratio of roughly 84%.[8]

This back-of-envelope calculation provides a rough but informative sense of the aggregate spillover magnitude. It is not a precise welfare measure: it does not net out intermediate transactions (so some double counting may arise between upstream and downstream losses), covers only first-degree propagation, abstracts from general equilibrium effects, and omits the cross-directional spillovers implied by the complementarity patterns documented in Section 7. The structural model developed in Section 8 provides a more rigorous quantification that addresses these limitations.

7 Mechanisms

The heterogeneity results above establish that firm size is the dominant driver of differential propagation effects. In this section, I investigate why large firms are better buffered. The core finding is that the degree of complementarity across trade relationships varies systematically with firm size. Small firms treat their trading partners as gross complements: losing one partner forces them to scale back trade with all other partners, amplifying the initial shock. Large firms, in contrast, behave as if their partners are gross substitutes: losing one partner has little effect on—or even increases—trade with remaining partners. This size-varying complementarity operates on both the input side (down-

---

[7] An alternative approach would be to estimate a sales-weighted regression directly. However, firm size is strongly correlated with the exposure measure (trade share to Xiong'an), so weighting by size would systematically down-weight the high-exposure observations that identify the propagation effect, introducing substantial bias. Estimating unweighted coefficients separately by firm size and then aggregating with appropriate sales weights avoids this problem.

[8] The direct loss is computed as a 53% decline (Table 3, Column 4: $1 - e^{-0.755} = 53.0\%$) applied to Xiong'an firms' total pre-policy sales (1,381 firms × 10.71M RMB mean sales = 14.8B RMB), yielding a direct loss of approximately 7.8B RMB per semi-annual period. The upstream small-firm spillover loss and downstream small-firm spillover loss are computed by applying the size-specific propagation coefficients (Table 6), adjusted for small firms' higher average trade exposure, to total small-firm sales. See Appendix C for details.

stream firms losing a supplier reduce purchases from other suppliers) and the output side (upstream firms losing a customer reduce sales to other customers). While input complementarity is a familiar channel in the production networks literature, the presence of strong output complementarity—and its sharp variation with firm size—is a novel empirical finding of this paper.

I present three pieces of evidence organized as follows. Section 7.1 documents the core result: size-varying complementarity on both the input and output sides. This is the central mechanism through which firm size shapes shock propagation. Section 7.2 then asks why complementarity varies with size, showing that the number of pre-existing alternative partners—which is strongly correlated with firm size—buffers shocks, while potential market partners provide no protection. Section 7.3 examines whether firms can offset the shock by forming new connections, finding that new link formation is limited overall, reinforcing the conclusion that shock propagation is governed by the complementarity structure of pre-existing networks rather than dynamic adjustment at the extensive margin.

7.1 Input and Output Complementarity

I begin with the central mechanism: size-varying complementarity across trade relationships. The key test is whether losing one trading partner causes a firm to reduce trade with its other, unaffected partners. If it does, the firm's partners are gross complements; if it does not, they are gross substitutes. I decompose firm-level outcomes into (a) trade with affected partners and (b) trade with other, unaffected partners, and estimate the propagation specification (Equation 2) separately for each component $k \in \{AFF, OTH\}$:[9]

$$y_{isct}^{k} = \sum_{\tau \neq -1} \beta_{\tau}^{k} \cdot \text{Exp\_XA}_{i}^{(1)} \times \text{Time}_{\tau} + \delta_{i} + \gamma_{sct} + \mathbf{X}'_{i,\tau_0 - 1}\theta_{t} + \varepsilon_{isct}, \quad k \in \{AFF, OTH\} \tag{4}$$

The coefficient β_{τ}^{AFF} captures the mechanical reduction in trade with the affected partner. The coefficient β_{τ}^{OTH} captures complementarity: a negative and significant β_{τ}^{OTH} indicates that firms that lose a key partner also scale back trade with their remaining partners.

To see why complementarity might arise, consider the supply chain from Figure 2. For the downstream auto manufacturer (input complementarity), the intuition is straightforward. When its Xiong'an steel supplier shuts down, the auto manufacturer cannot simply use more of its other inputs—production requires a specific mix of steel, rubber, glass, and labor, and reducing one constrains the use of others. Without a quick replacement, the firm scales back overall production, reducing purchases from all its suppliers. This is the classic input complementarity channel emphasized in the production networks literature (Acemoglu et al., 2012; Carvalho et al., 2021; Boehm et al.,

---

[9] For upstream firms, "other customers" excludes customers of Xiong'an industrial firms, to rule out confounding supply shocks among overlapping customers. An analogous exclusion applies to downstream firms.

2019).

For the upstream ore supplier (output complementarity), the mechanism is less obvious but equally important. When the steel mill shuts down, why would the ore supplier also reduce sales to its other customers? The answer lies in scale economies: a firm that loses a major customer sees its effective scale shrink, which erodes productive capacity and limits its ability to serve remaining buyers.[10] The structural model in Section 8 formalizes this as demand-dependent TFP, parameterized by φ_i, without taking a stand on the particular micro channel. Small firms—with thinner financial buffers, less diversified operations, and fewer alternative outlets—should exhibit stronger output complementarity.

Figure 9 plots the event-study coefficients from the decomposed specification. Panels (a) and (c) confirm the mechanical effect: more exposed firms reduce trade with their affected partners (Xiong'an firms). Critically, Panels (b) and (d) reveal significant complementarity: upstream firms with higher exposure also reduce sales to their other (unaffected) customers, and downstream firms reduce purchases from their other suppliers.

Table 7 decomposes these complementarity effects by firm size. The results are striking and confirm the size-varying complementarity mechanism. For upstream firms selling to other customers (excluding customers of Xiong'an firms), small firms show a significant 9.7% reduction in sales—strong evidence of *output complementarity*. The size interaction, while positive (8.3 percentage points), substantially attenuates this effect: the implied net effect on large firms is only −1.4%, which is statistically insignificant, indicating that large firms' customers are effectively gross substitutes. The pattern is even cleaner for downstream firms purchasing from other suppliers (excluding suppliers of Xiong'an firms): small firms show a significant 6.4% reduction in purchases—evidence of *input complementarity*—while the size interaction of 10.6 percentage points fully offsets and reverses this baseline. The implied net effect for large firms is a *positive* 4.3%, meaning they actually increase purchases from other suppliers, consistent with active substitution toward remaining partners. These results reveal a sharp dichotomy: for small firms, trade partners are gross complements on both the input and output sides, so that losing one partner cascades into broader contraction. For large firms, trade partners are effectively gross substitutes, allowing them to absorb the shock without disrupting their other relationships.

---

[10] Several channels can generate this linkage: reduced cash flow constraining production capacity (financial constraints, binding most tightly for small firms); higher average costs from spreading fixed infrastructure across fewer customers; degradation of customer-specific capabilities that spill over across the customer base; and dedicated capacity that cannot be easily redirected to buyers with different specifications. All share a common reduced-form prediction: losing a customer contracts the firm's effective productivity, and this contraction spills over to remaining relationships.

Figure 9: Event Study Results: Decomposition of the Propagation Effects

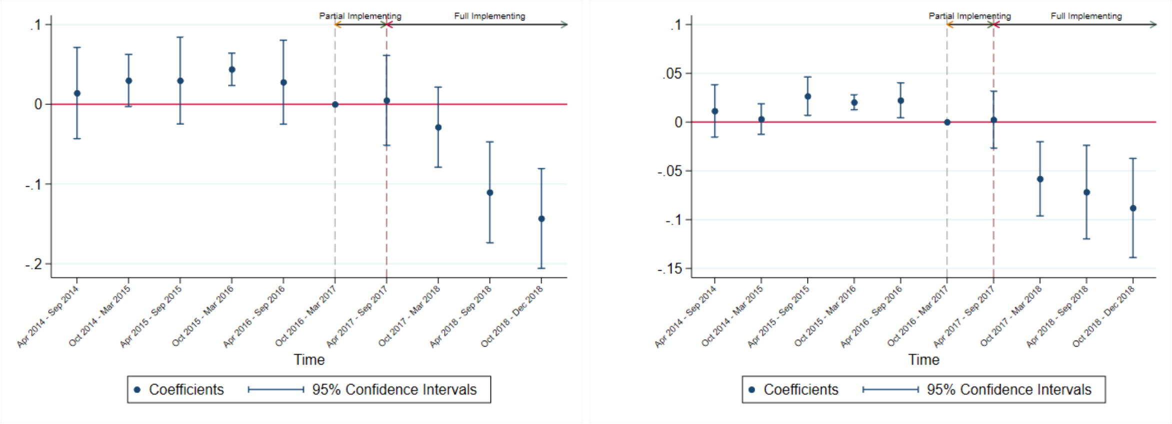

(a) Upstream (Log Sales to Affected Customers)

(b) Upstream (Log Sales to Other Customers)

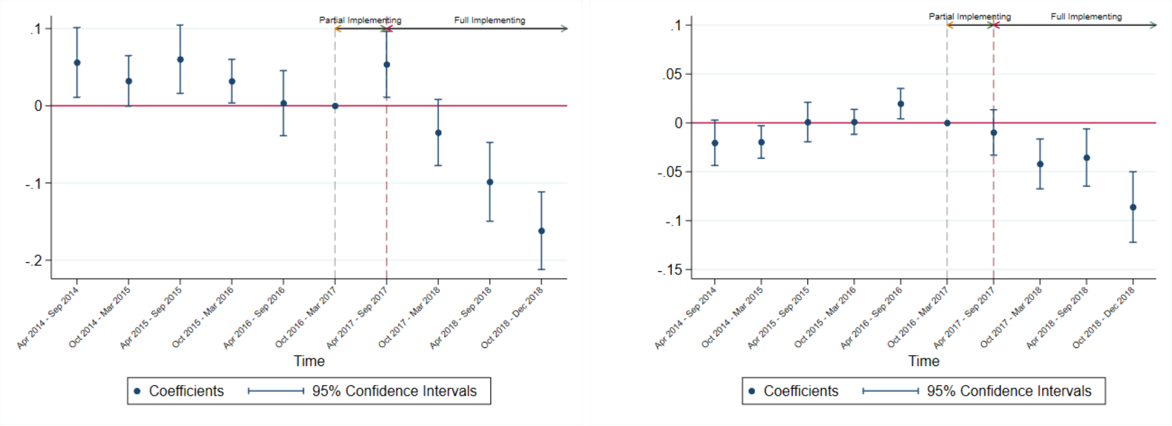

(c) Downstream (Log Purchases from Affected Suppliers)

(d) Downstream (Log Purchases from Other Suppliers)

Notes: This figure plots event-study coefficients from the propagation specification (Equation 2) using decomposed outcome variables. Panels (a) and (b) use the upstream sample: (a) uses log sales to affected customers (Xiong'an firms for firms with positive exposure; pseudo-affected matched control firms' customers for firms with zero exposure) as the outcome; (b) uses log sales to other customers, excluding customers of Xiong'an industrial firms to rule out confounding supply shocks among overlapping customers. Panels (c) and (d) use the downstream sample: (c) uses log purchases from affected suppliers; (d) uses log purchases from other suppliers, excluding suppliers of Xiong'an firms to avoid contamination from shared demand shocks. The decline in Panels (b) and (d) indicates complementarity—firms reduce trade with unaffected partners when they lose an affected partner. The omitted reference period is October 2016–March 2017 ($\tau = -1$). Standard errors are clustered at the firm level.

Table 7: Static DID Results: Decomposition of the Propagation Effects by Firm Size

	Sample: 1st degree upstream firms		Sample: 1st degree downstream firms	
	log_sale_affected	log_sale_other	log_purchase_affected	log_purchase_other
	(1)	(2)	(3)	(4)
$\text{UpExp}^{(1)} \times \text{Full}$	-0.1228***	-0.0972***		
	(0.0268)	(0.0231)		
$\text{UpExp}^{(1)} \times \text{Full} \times \text{Large}$	0.0356	0.0834		
	(0.0745)	(0.0520)		
$\text{DownExp}^{(1)} \times \text{Full}$			-0.1349***	-0.0635***
			(0.0234)	(0.0152)
$\text{DownExp}^{(1)} \times \text{Full} \times \text{Large}$			0.0524	0.1061***
			(0.0631)	(0.0288)
Observations	1,499,860	1,499,860	1,757,040	1,757,040
R^2	0.5707	0.7308	0.5707	0.8006
Firm FE	X	X	X	X
Time-Province-Industry FE	X	X	X	X
Pre-Policy Controls × Time FE	X	X	X	X
Implied effect on large firms	-0.0872	-0.0138	-0.0825	0.0426*
	(0.0699)	(0.0464)	(0.0593)	(0.0249)
Mean 1st degree % sales to XiongAn	0.05	0.05		
Mean 1st degree % purchases from XiongAn			0.07	0.07

Notes: This table presents the decomposition of propagation effects by firm size. Columns (1)–(2) use the upstream sample with decomposed sales outcomes; Columns (3)–(4) use the downstream sample with decomposed purchase outcomes. "Sale Affected" ("Purchase Affected") refers to sales to (purchases from) affected customers (suppliers)—i.e., Xiong'an firms or their pseudo-affected matched controls. "Sale Other" ("Purchase Other") refers to trade with all other partners, excluding customers (suppliers) of Xiong'an industrial firms to rule out confounding shocks from overlapping network connections. "Large" is an indicator for above-median firm size. The negative baseline coefficients on trade with "Other" partners indicate complementarity: losing an affected partner causes firms to reduce trade with unaffected partners. The positive interaction with "Large" shows that this complementarity is concentrated among small firms. For large firms, the implied effect on "Other" trade is small and statistically insignificant for upstream (−1.4%) and significantly positive for downstream (+4.3%), indicating that large firms treat their partners as gross substitutes and can reallocate trade accordingly. All specifications include firm fixed effects, industry-province-time fixed effects, and pre-policy firm characteristics interacted with time fixed effects. Standard errors clustered at the firm level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

7.2 Pre-Existing Alternative Partners

The complementarity results above raise a natural question: why does complementarity vary so sharply with firm size? A direct explanation is that large firms have more pre-existing alternative partners for the same products, giving them the capacity to substitute away from a lost partner without scaling back overall operations. I now provide direct evidence for this channel.

I leverage the detailed product-level information that became available in the VAT data starting in 2017. The State Taxation Administration mandated a new commodity and service classification system, a comprehensive domestic equivalent to the international Harmonized System (HS). I define distinct products using the 7-digit commodity codes (roughly equivalent to 4-digit HS codes) to capture broader product categories where reallocation might be more feasible. My data contain approximately 800 unique 7-digit codes. A distinct advantage of these VAT codes over traditional HS codes is their coverage of both commodities and services (e.g., logistics, consulting), allowing me to capture the full spectrum of input-output linkages. For brevity, I refer to both goods and services as "products" throughout this section.

I construct a measure of pre-existing alternative connections based on a firm's pre-policy relationships for the specific products it traded with Xiong'an. For an upstream supplier i, I first identify the set of products $P_{i\mathcal{S}}$ it sold to shut-down firms in $\mathcal{S}$. For each product $p \in P_{i\mathcal{S}}$, I count the number of existing alternative customers (N_{ip}^{alt}) for that same product. My measure of alternative connections, S_i^{ALT}, is the weighted average of these counts, where weights w_{ip} are based on the product's share of firm i's total sales to $\mathcal{S}$: $S_i^{ALT} = \sum_{p \in P_{i\mathcal{S}}} w_{ip} \cdot N_{ip}^{alt}$. A larger value means the supplier already has established relationships with other buyers for the same products. I construct analogous measures for downstream customers based on their pre-existing alternative suppliers.

Table 8 shows the results. For upstream suppliers, the baseline effect is an 11.9% sales drop for mean-exposed suppliers, but having one additional pre-existing customer for the affected products mitigates this shock by 0.26 percentage points. Given that suppliers in my sample have, on average, 119 existing alternative partners, this buffering effect can be economically substantial. For downstream customers, I find a similar buffering effect for purchases: having one additional alternative supplier helps mitigate the purchase contraction by 1.9 percentage points.

Critically, I document a strong correlation between alternative connections and firm size. As shown in Figure 10, larger firms tend to have substantially more existing alternative trade partners. This correlation provides the micro-foundation for the size-varying complementarity documented in Section 7.1: large firms show no complementarity precisely because they have enough pre-existing alternatives to substitute away from a lost partner. Small firms, with fewer alternatives, cannot substitute and are forced to contract—hence the gross complementarity in their trade relationships.

I additionally examined whether the availability of potential new partners in the broader market (i.e., all firms trading the same products but not currently connected) provides any buffering. Despite vast potential markets (average: 308,000 potential partners for upstream, 130,000 for downstream), I find no buffering effect—the interaction coefficients are essentially zero and statistically insignificant (see Appendix Table B.5). This null result reinforces the importance of pre-existing relationships: the supply chain is highly "sticky," and the mere existence of market liquidity does not help firms replace lost partners.

As a robustness check, I repeated this analysis using the highly granular 13-digit product codes (roughly equivalent to 8-digit HS codes). The results, reported in Appendix Table B.6 and Table B.7, are qualitatively similar to those using 7-digit codes, confirming that my findings are not driven by product aggregation.

Figure 10: Correlation between Firm Size and Pre-Existing Alternative Partners

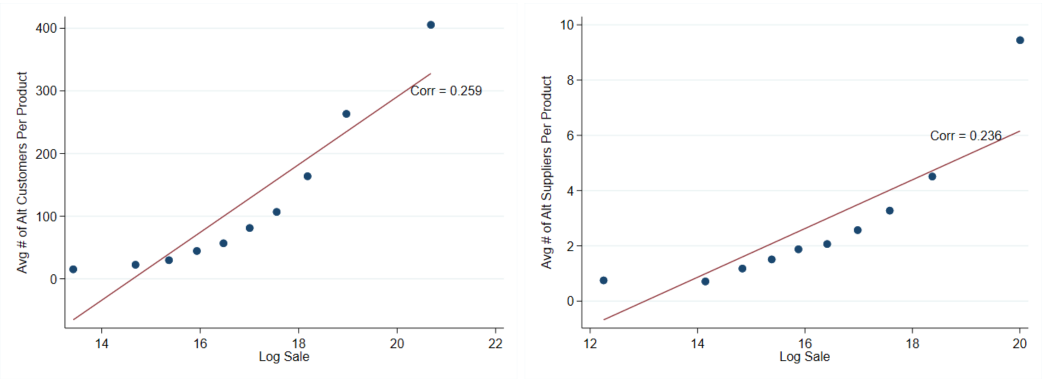

(a) Upstream (b) Downstream

Notes: This figure plots binned scatterplots of firm size (log pre-policy total sales, x-axis) against the number of pre-existing alternative partners (S^{ALT}, y-axis) for upstream suppliers (Panel a) and downstream customers (Panel b). The positive relationship confirms that larger firms have substantially more pre-existing alternative partners for the specific products traded with Xiong'an, explaining why firm size buffers propagation effects.

Table 8: Heterogeneous Propagation Effects: Pre-Existing Alternative Partners (7-digit Codes)

	Sample: 1st degree upstream firms		Sample: 1st degree downstream firms	
	log_purchase	log_sale	log_purchase	log_sale
	(1)	(2)	(3)	(4)
$\text{UpExp}^{(1)} \times \text{Full}$	-0.0642***	-0.1188***		
	(0.0174)	(0.0221)		
$\text{UpExp}^{(1)} \times \text{Full} \times S^{INT}$	0.0015***	0.0026***		
	(0.0004)	(0.0007)		
$\text{DownExp}^{(1)} \times \text{Full}$			-0.1032***	-0.0865***
			(0.0159)	(0.0179)
$\text{DownExp}^{(1)} \times \text{Full} \times S^{INT}$			0.0190***	-0.0024
			(0.0071)	(0.0128)
Observations	1,474,430	1,474,430	1,714,400	1,714,400
R^2	0.7965	0.7244	0.8069	0.6617
Firm FE	X	X	X	X
Time-Province-Industry FE	X	X	X	X
Pre-Policy Controls × Time FE	X	X	X	X
Mean 1st degree % sales to XiongAn	0.08	0.08		
Mean 1st degree % purchases from XiongAn			0.12	0.12
Avg. S^{INT} (Existing Partners)	118.87	118.87	2.79	2.79

Notes: This table presents heterogeneous effects of the first-degree propagation by the number of pre-existing alternative partners (S^{ALT}), defined using 7-digit product codes. S^{ALT} measures the firm's weighted average number of pre-existing alternative partners for the products traded with Xiong'an. The interaction coefficient represents the differential effect of having one additional pre-existing partner. Note that S^{ALT} is strongly correlated with firm size: larger firms have substantially more pre-existing alternative partners on average (Avg. S^{ALT}: 118.87 for upstream, 2.79 for downstream). All specifications include firm fixed effects, industry-province-time fixed effects, and pre-policy firm characteristics interacted with time fixed effects. Standard errors clustered at the firm level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

7.3 New Connection Formation

A remaining question is whether firms can mitigate the shock by forming new trade connections after the disruption. If firms could readily establish new relationships, the complementarity patterns documented above would be attenuated over time as firms replace lost partners. I now show that this is not the case.

To examine extensive margin responses, I construct two measures of new connection formation. Using the period October 2014–March 2015 as the "legacy" baseline, I identify new partners in each subsequent period t as those that never appeared in any prior period. Let $\mathcal{N}_{i,t}^{new}$ denote firm i's set of new partners in period t and $\mathcal{N}_{i,\text{pre}}$ denote its pre-policy partner set. The first measure, the new partner rate, $|\mathcal{N}_{i,t}^{new}| \,/\, |\mathcal{N}_{i,\text{pre}}|$, captures the count of new partners relative to the firm's pre-shock partner count. The second, the new trade share, $\sum_{j\in\mathcal{N}_{i,t}^{new}} \text{Trade}_{ij,t} \,/\, \sum_{j\in\mathcal{N}_{i,\text{pre}}} \text{Trade}_{ij,\text{pre}}$, captures the trade value with new partners relative to pre-shock trade volume. Both denominators are fixed at pre-period levels to avoid the endogeneity that arises when the contemporaneous partner count responds to the shock.[11] I estimate the effect of Xiong'an exposure on these outcomes using the same propagation specification as Equation 2, with the new connection measures as the dependent variables.

Appendix Figure A.4 plots the event-study coefficients. On the upstream side, more exposed suppliers show no clear increase in new customer formation—the post-shock coefficients are small and imprecise. On the downstream side, more exposed customers show modest increases in new supplier formation, but the magnitudes are small relative to pre-period base rates (5% for new customer rate, 14% for new supplier rate). Table 9 further shows that size heterogeneity in new link formation is limited: neither small nor large upstream firms show significant increases in new customer formation, and while small downstream firms show modest increases in new supplier formation, the size interactions are largely insignificant. Even in the face of a severe, permanent shock, new link formation rates remain low in absolute terms, reinforcing the network stickiness documented in Section 7.2 and confirming that shock propagation is governed by the complementarity structure of pre-existing networks.

---

[11] With a contemporaneous denominator, a firm that loses partners due to the shock would mechanically show a higher new-partner rate even with no actual change in new link formation, biasing the estimates upward.

Table 9: Static DID Results: New Connection Formation by Firm Size

	Sample: 1st degree upstream firms		Sample: 1st degree downstream firms	
	rate_new_customer	rate_new_sale	rate_new_supplier	rate_new_purchase
	(1)	(2)	(3)	(4)
$\text{UpExp}^{(1)} \times \text{Full}$	-0.0007	0.0014		
	(0.0011)	(0.0011)		
$\text{UpExp}^{(1)} \times \text{Full} \times \text{Large}$	-0.0025	-0.0057**		
	(0.0026)	(0.0023)		
$\text{DownExp}^{(1)} \times \text{Full}$			0.0020**	0.0061***
			(0.0010)	(0.0009)
$\text{DownExp}^{(1)} \times \text{Full} \times \text{Large}$			0.0009	-0.0038*
			(0.0020)	(0.0022)
Observations	1,207,880	1,207,800	1,421,824	1,421,336
R^2	0.6263	0.4728	0.6100	0.3728
Firm FE	X	X	X	X
Time-Province-Industry FE	X	X	X	X
Pre-Policy Controls × Time FE	X	X	X	X
Implied effect on large firms	-0.0032	-0.0043**	0.0029*	0.0023
	(0.0024)	(0.0021)	(0.0018)	(0.0020)
Pre-period Mean outcome value	0.05	0.04	0.14	0.07
Mean 1st degree % sales to XiongAn	0.05	0.05		
Mean 1st degree % purchases from XiongAn			0.07	0.07

Notes: This table presents the effects of Xiong'an exposure on new connection formation, separately for small and large firms. Columns (1)–(2) use the upstream sample; Columns (3)–(4) use the downstream sample. "Rate New Customer/Supplier" is the number of new partners divided by the pre-period total partner count. "Rate New Sale/Purchase" is the value of trade with new partners divided by pre-period total trade volume. The denominators are fixed at pre-period levels to avoid mechanical inflation from contemporaneous partner loss. "Large" is an indicator for above-median pre-policy firm size. New partners are defined as those that never appeared prior to October 2014–March 2015. The results show limited and mixed evidence on new connection formation: upstream firms show no significant increase in new customer formation regardless of size, while downstream firms show modest increases that do not differ strongly by firm size. Overall, the low absolute rates of new link formation—relative to pre-period base rates of 5% (upstream) and 14% (downstream)—confirm substantial network stickiness. All specifications include firm fixed effects, industry-province-time fixed effects, and pre-policy firm characteristics interacted with time fixed effects. Standard errors clustered at the firm level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

7.4 Summary: The Mechanism

To summarize, the mechanism behind heterogeneous shock propagation is size-varying complementarity across trade relationships. Small firms treat their partners as gross complements on both the input and output sides: losing one partner forces contraction with all others. Large firms treat their partners as gross substitutes, absorbing partner loss without disrupting remaining relationships. This complementarity varies with size because larger firms have substantially more pre-existing alternative partners, giving them the capacity to substitute; the mere existence of potential market partners provides no buffering. New connection formation is limited overall, confirming that firms cannot escape this complementarity structure through dynamic adjustment. Together, these findings establish a pattern I term "Network Regressivity": small firms amplify shocks through complementarity-driven contraction while large firms absorb them through substitution, so that supply chain disruptions disproportionately burden the smallest and most peripheral firms. The structural model below formalizes this mechanism and derives its welfare implications.

8 Model

The reduced-form analysis establishes that regulatory shocks propagate through production networks and that the resulting spillover costs fall predominantly on small firms, driven by two-sided complementarity across their trade relationships. A structural model complements these findings in three ways. First, it maps the reduced-form elasticities into an aggregate GDP loss by assigning each firm a general equilibrium weight (Domar weight), enabling a channel decomposition that separates the contributions of input substitution (ξ), demand-dependent TFP (φ), and factor reallocation (σ). Second, it delivers a welfare criterion that equates the social marginal cost of abatement—including network spillovers—with the marginal benefit of emission reductions, pinning down the socially optimal regulation level. Third, the framework supports counterfactual comparisons between policy instruments, in particular between the observed quantity control (the Xiong'an shutdown) and an equivalent emission tax.

8.1 Model Setup

I build on the production network framework of Carvalho et al. (2021), extending it with two innovations: demand-dependent TFP (scale economies) and size-varying elasticities.

Production Technology. Consider an economy with N competitive firms. Each firm i produces output using labor l_i and a bundle of intermediate inputs M_i:

$$y_i = z_i^{eff} \left[\mu^{1/\sigma} M_i^{(\sigma-1)/\sigma} + (1-\mu)^{1/\sigma} l_i^{(\sigma-1)/\sigma} \right]^{\sigma/(\sigma-1)} \tag{5}$$

where μ is the intermediate input share and σ is the elasticity of substitution between intermediates and labor. The key departure from standard models is that the effective TFP z_i^{eff} depends on downstream conditions, as described below.

The intermediate input bundle M_i is a CES aggregate of goods purchased from other firms:

$$M_i = \left(\sum_{j=1}^{N} a_{ij}^{1/\xi_i} x_{ij}^{(\xi_i-1)/\xi_i} \right)^{\xi_i/(\xi_i-1)} \tag{6}$$

where x_{ij} is the quantity of good j used by firm i, a_{ij} is the taste parameter (calibrated to pre-shock expenditure shares), and ξ_i is the elasticity of substitution across intermediate inputs.

Demand-Dependent TFP (Scale Economies). Standard production network models assume firm TFP depends only on exogenous technology z_i. I allow effective TFP to depend on downstream conditions: when a firm's customers contract, its productive capacity erodes. In log-changes, effective TFP satisfies:

$$\theta_i^{eff} \equiv d\log z_i^{eff} = \underbrace{d\log z_i}_{\theta_i \text{ (exogenous)}} + \varphi_i \underbrace{\sum_k \tilde{a}_{ik}(L\boldsymbol{\theta}^{eff})_k}_{G_i(\boldsymbol{\theta}^{eff}) \text{ (customer health)}} \tag{7}$$

where $\tilde{a}_{ik}$ is firm i's pre-shock sales share to customer k, $L = (I - \mu A)^{-1}$ is the standard Leontief inverse, and $\varphi_i \geq 0$ governs the strength of this demand–productivity linkage.

This formulation captures the output complementarity documented in Section 7.1: small upstream firms that lose Xiong'an buyers also reduce sales to their unaffected customers, consistent with demand-dependent TFP operating through the scale economy channels discussed there. The parameter φ_i is estimated directly from the data without taking a stand on the particular micro channel generating the demand–productivity linkage.

The customer base is not a choice variable for the firm—it is determined by downstream demand conditions. From the firm's optimization perspective, z_i^{eff} is a parameter, so the production function exhibits constant returns to scale in the choice variables (M_i, l_i), preserving competitive equilibrium.[12]

Rather than modeling the demand-side feedback through explicit demand functions and a separate externality $\mathcal{D}_i$, I absorb it directly into effective TFP. This does not sacrifice generality. For competitive firms, a demand externality entering multiplicatively and a direct TFP formulation $z_i^{eff} = z_i \cdot \mathcal{D}_i$ yield identical equilibrium prices, quantities, Domar weights, and GDP up to second order (Appendix D). This follows a general principle in efficient production network economies: the source of a shock is irrelevant for equilibrium outcomes when the affected agent does not internalize it. Baqaee and Farhi (2019) establish a related result—that factor supply shocks can always be restated as productivity shocks—and Huo et al. (2025) apply this in the international context.[13] The z_i^{eff} formulation achieves the same economic content while preserving the tractable sequential solution structure described below.

---

[12] The CES production technology is homogeneous of degree one in (M_i, l_i) for any value of the substitution elasticities σ and ξ_i. Since z_i^{eff} enters multiplicatively and is taken as given, it does not affect the degree of homogeneity. Average cost equals marginal cost at all output levels, sustaining price-taking behavior. See Appendix D for a formal discussion.

[13] The BF result concerns factor supply ↔ TFP isomorphism; my normalization concerns demand externality ↔ TFP isomorphism, which relies on a different but analogous argument: the firm's FOCs depend only on $z_i^{eff} = z_i \cdot \mathcal{D}_i$, not on whether $\mathcal{D}_i$ originates from intrinsic productivity or downstream conditions. The common logic is that in competitive equilibrium with CRS, effective quantities are sufficient statistics regardless of source. Under monopolistic competition with variable markups, the equivalence would partially break down because demand shifts and TFP shifts affect the markup through different elasticities. Under CES demand (constant markups), the distinction vanishes. The competitive framework used here avoids the issue entirely; see Appendix D for a formal proof.

Size-Varying Parameters. Both substitution and scale parameters vary with firm size:

$$\xi_i = \xi_0 + \xi_1 \cdot \mathbf{1}\{\text{Large}_i\} \tag{8}$$

$$\varphi_i = \varphi_0 + \varphi_1 \cdot \mathbf{1}\{\text{Large}_i\} \tag{9}$$

The mechanism analysis predicts $\xi_1 > 0$ (large firms have higher input substitutability, as they can reallocate across more pre-existing alternative partners) and $\varphi_1 < 0$ (large firms have weaker scale sensitivity, as they are less dependent on any single customer).

Household. A representative household has Cobb-Douglas preferences with expenditure shares $\{\beta_i\}$:

$$p_i c_i = \beta_i Y \quad \text{for all } i \tag{10}$$

Equilibrium. A competitive equilibrium is a set of prices $\{p_i\}$, wage w, and quantities $\{y_i, x_{ij}, c_i, l_i\}$ satisfying household optimality, firm cost minimization, and market clearing ($y_i = c_i + \sum_k x_{ik}$ for all i). Because the firm treats z_i^{eff} as a parameter, the equilibrium conditions are standard—conditional on the effective TFP vector, prices and quantities are determined by the usual Leontief input–output structure. The only nonstandard step is tracking how z_i^{eff} responds to shocks through the fixed-point condition (7).

8.2 Shock Transmission

I characterize shock transmission in four steps, following the perturbation approach of Baqaee and Farhi (2019) augmented with the scale channel.

Step 1: Scale Amplification. Before any network propagation through the Leontief inverse, the scale channel amplifies exogenous shocks $\boldsymbol{\theta}$ into effective shocks via a fixed-point condition. In matrix form, Equation (7) becomes:

$$\boldsymbol{\theta}^{eff} = \boldsymbol{\theta} + \text{diag}(\boldsymbol{\varphi})\, \tilde{A}L\, \boldsymbol{\theta}^{eff} \tag{11}$$

Rearranging:

$$\boldsymbol{\theta}^{eff} = \underbrace{(I - \text{diag}(\boldsymbol{\varphi})\, \tilde{A}L)^{-1}}_{\equiv L^{SE} \text{ (scale-extended Leontief)}} \boldsymbol{\theta} \tag{12}$$

The operator L^{SE} has a convergent Neumann series when $\max_i(\varphi_i) < 1/\rho(\tilde{A}L)$:

$$L^{SE} = I + \text{diag}(\boldsymbol{\varphi})\tilde{A}L + [\text{diag}(\boldsymbol{\varphi})\tilde{A}L]^2 + \cdots \tag{13}$$

Each successive power adds another round of customer-side feedback. The zeroth-order term (I) gives the exogenous shock; the first-order term gives the direct scale feedback; higher-order terms capture cascading scale amplification. When $\varphi = 0$, $L^{SE} = I$ and $\boldsymbol{\theta}^{eff} = \boldsymbol{\theta}$: the model reduces to the standard Carvalho et al. (2021) framework, and the remaining derivation follows their analytical machinery exactly. When $\varphi > 0$, all subsequent steps proceed identically but with $\boldsymbol{\theta}^{eff}$ replacing $\boldsymbol{\theta}$.

Step 2: Price Propagation. Given effective TFP shocks, prices propagate through the input network via the standard Leontief channel: $d \log p_i = -(L\boldsymbol{\theta}^{eff})_i$.

Step 3: Input Reallocation. When input prices change, firms reallocate purchases, creating second-order effects on Domar weights:

$$\Delta \log \lambda_i = (\sigma - 1)\Sigma_i(\boldsymbol{\theta}^{eff}) + (\xi_i - 1)\Xi_i(\boldsymbol{\theta}^{eff}) \tag{14}$$

The network statistics Σ_i and Ξ_i measure how unevenly price changes fall across a firm's inputs. Specifically, Σ_i captures dispersion across factor types (intermediates versus labor):

$$\Sigma_i = \mu_i(1 - \mu_i)\left(d \log P_i^M - d \log w\right)^2 \tag{15}$$

where μ_i is the intermediate input share and $d \log P_i^M - d \log w$ is the relative price change between intermediates and labor. When $\sigma > 1$, firms substitute toward the cheaper factor, expanding in relative terms and increasing their Domar weight; when $\sigma < 1$, the reverse occurs. The term Ξ_i captures dispersion across intermediate suppliers:

$$\Xi_i = \sum_j \omega_{ij}\left(d \log p_j - d \log P_i^M\right)^2 \tag{16}$$

where ω_{ij} is firm i's expenditure share on supplier j and $d \log p_j - d \log P_i^M$ is the deviation of supplier j's price change from the expenditure-weighted average. When $\xi_i > 1$, firms substitute toward cheaper suppliers, buffering losses; when $\xi_i < 1$, firms cannot substitute and losses are amplified. Both statistics are evaluated at the effective shock $\boldsymbol{\theta}^{eff}$. See Appendix D.3 for the full derivation.

Step 4: GDP Loss. Combining these channels, the second-order approximation to aggregate GDP loss is:

$$\boxed{\Delta \log Y = \underbrace{\boldsymbol{\lambda}^\top \boldsymbol{\theta}}_{\substack{\text{(a) Direct loss}\\ \text{(Hulten)}}} + \underbrace{\boldsymbol{\lambda}^\top \boldsymbol{\delta}}_{\substack{\text{(b) Scale}\\ \text{amplification}}} + \frac{1}{2}\underbrace{\sum_j (\lambda_j^* - \lambda_j^0)\theta_j^{eff}}_{\substack{\text{(c) Domar}\\ \text{reallocation}}} + \underbrace{\frac{1}{2}\mu(\sigma - 1)(\boldsymbol{\theta}^{eff})^\top \boldsymbol{\Lambda}(\mathbf{I} - \mathbf{A})\mathbf{L}\,\boldsymbol{\theta}^{eff}}_{\substack{\text{(d) Intermediate share}\\ \text{adjustment}}}} \tag{17}$$

where $\boldsymbol{\delta} \equiv \boldsymbol{\theta}^{eff} - \boldsymbol{\theta}$ is the scale-induced amplification ($\boldsymbol{\delta} = \mathbf{0}$ when $\varphi = 0$) and $\lambda_j^* = \lambda_j^0(1 + (\sigma - 1)\Sigma_j + (\xi_j - 1)\Xi_j)$ is the post-shock Domar weight.

Channel (a) is the standard first-order Hulten term, using only the exogenous shock $\boldsymbol{\theta}$ and pre-shock Domar weights. Channel (b) is the additional first-order loss from demand-dependent TFP—this channel is new relative to standard production network models and is proportional to φ. Channel (c) captures how market shares shift toward or away from affected firms, depending on the substitution elasticities σ and ξ_i. Channel (d) captures how the economy-wide intermediate-to-value-added ratio responds when $\sigma \neq 1$, following the derivation in Carvalho et al. (2021).

The sequential solution—computing $\boldsymbol{\theta}^{eff}$ at pre-shock weights before evaluating substitution effects—introduces an approximation. The error can be characterized using the chain rule for $\log Y = f(\boldsymbol{\theta}^{eff}(\boldsymbol{\theta}))$: the second-order GDP cross-partial decomposes into Source A (standard Baqaee–Farhi reallocation, which both procedures agree on) and Source B (curvature of the scale mapping, which the sequential procedure sets to zero). Source B scales as $\varphi \cdot \varepsilon \cdot \theta^2$ where $\varepsilon = \max\{|\sigma - 1|, |\xi - 1|\}$—formally $O(\theta^2)$ but attenuated by φ relative to the kept second-order terms. For plausible parameters, Source B is roughly 0.6% of total GDP loss. The error vanishes when $\varphi = 0$ or $\sigma = \xi = 1$. The practical resolution is to estimate parameters sequentially, then validate with an iterated counterfactual that solves the simultaneous equilibrium at the estimated $\hat{\boldsymbol{\psi}}$ (Appendix D.2).[14]

8.3 Estimation Strategy: Minimum Distance

I estimate the parameter vector $\boldsymbol{\psi} = (\sigma, \xi_0, \xi_1, \varphi_0, \varphi_1)$ using a minimum-distance approach that matches model predictions to within-firm ratio moments constructed from the same subsamples and control groups as the reduced-form analysis. The key advantage of this approach is a block-diagonal identification structure: each parameter is identified from its own set of moments, with no cross-contamination.

I target $K = 6$ empirical moments organized in three blocks.

Moments 1–2: Purchase-to-Sales Ratio (Identifying σ). Cost minimization implies that the materials expenditure share $s_i^M \equiv P_i^M M_i/(p_i y_i)$ depends on the relative price of intermediates and labor through σ. Under perfect competition, $\text{Purchases}_i/\text{Sales}_i = s_i^M$. Log-differentiating:

$$\Delta \log \frac{\text{Purch}_i}{\text{Sales}_i} = (1 - \sigma) \cdot (1 - \mu) \cdot \Big(d \log P_i^M - d \log w\Big) \tag{18}$$

In a DID between treated (exposed) and control firms, the wage change is absorbed by sector×county×time fixed effects, leaving the materials price change as the sole regressor. Since

---

[14] Carvalho et al. (2021) follow the same two-stage approach: they estimate from a linearized model (their Table III) and compute counterfactual GDP losses from the exact nonlinear equilibrium (their Table V). My setting adds the φ fixed point to the iteration but the logic is identical.

σ is a technology parameter common across firm sizes, I pool small and large firms within each position group, yielding two moments: $\hat{m}_1$ (pooled downstream) and $\hat{m}_2$ (pooled upstream). Both identify the same $(1-\sigma)(1-\mu)$, providing one overidentifying restriction. The structural coefficient is invariant to φ: the scale channel affects the magnitude of the price change but not the elasticity of the purchase-to-sales ratio with respect to that price change.

Moments 3–4: Expenditure Ratio (Identifying ξ). Within the intermediates bundle, CES cost minimization gives the ratio of spending on affected versus unaffected suppliers:

$$\Delta \log \frac{\text{Purch}_i^{AFF}}{\text{Purch}_i^{OTH}} = (1-\xi_i) \cdot \underbrace{\left(\overline{\Delta \log p_i}^{AFF} - \overline{\Delta \log p_i}^{OTH}\right)}_{\equiv \Delta p_i^{GAP}} \tag{19}$$

where Δp_i^{GAP} is the within-firm price gap between affected and unaffected suppliers, computed using pre-shock (Laspeyres) expenditure weights. Total intermediate spending cancels in the ratio, removing any dependence on σ. This approach follows Boehm et al. (2019), who exploit within-firm variation in supplier exposure to identify input substitution elasticities. I estimate the moment in a DID regression of the log expenditure ratio on $\Delta p_i^{GAP} \times \text{Post}_t$, with a size interaction to identify ξ_1 separately from ξ_0. The slope recovers $(1-\xi_0)$ directly, and the size interaction recovers $-\xi_1$. See Appendix D.4 for the full regression specification.

Moments 5–6: Upstream Price DID (Identifying φ). The firm's own price identifies the scale channel. For an upstream firm not itself targeted ($\theta_i = 0$), effective TFP declines through customer health: $\theta_i^{eff} = \varphi_i G_i + O(\varphi^2)$, where $G_i = [\tilde{A}L\boldsymbol{\theta}]_i$ is the firm's sales-share-weighted average customer exposure.[15] Log-differentiating unit cost:

$$d \log p_i = -\theta_i^{eff} + s_i^M \, d \log P_i^M + s_i^L \, d \log w \tag{20}$$

Controlling for input price pass-through and absorbing the wage change with fixed effects, the residual price increase identifies φ: the DID coefficient on $G_i \times \text{Post}_t$ recovers $-\varphi_0$ at first order, and its size interaction recovers $-\varphi_1$. Upstream firms whose customers are more exposed to the shutdown exhibit larger price increases, reflecting the erosion of their effective TFP through the scale channel. See Appendix D.4 for the full regression specification.

The ratio-based identification for sales—natural for ξ on the input side—does not work for φ on the output side. Under single-good production, firm i sells at a uniform price p_i to all customers, so any sales ratio cancels p_i and is insensitive to φ. The price level is what reveals the demand–

---

[15] The model defines $G_i(\boldsymbol{\theta}^{eff}) = [\tilde{A}L\boldsymbol{\theta}^{eff}]_i$ using the effective shock, but in the estimation I use $G_i = [\tilde{A}L\boldsymbol{\theta}]_i$ with the exogenous shock. The difference is $O(\varphi)$, consistent with the $O(\varphi^2)$ approximation.

productivity linkage.

Minimum-Distance Estimator. The estimator minimizes:

$$\hat{\boldsymbol{\psi}} = \arg\min_{\boldsymbol{\psi}} \left[\hat{\mathbf{m}} - \mathbf{m}(\boldsymbol{\psi})\right]^\top \hat{\mathbf{V}}^{-1} \left[\hat{\mathbf{m}} - \mathbf{m}(\boldsymbol{\psi})\right] \tag{21}$$

where $\hat{\mathbf{m}} = (\hat{m}_1, \ldots, \hat{m}_6)^\top$ is the vector of empirical moments, $\mathbf{m}(\boldsymbol{\psi})$ is the vector of model predictions, and $\hat{\mathbf{V}}$ is the bootstrapped variance-covariance matrix. The Jacobian $\mathbf{B} = \partial \mathbf{m}/\partial \boldsymbol{\psi}$ is block-diagonal: moments 1–2 load on σ only, moments 3–4 load on (ξ_0, ξ_1) only, and moments 5–6 load on (φ_0, φ_1) only. Each parameter block can therefore be estimated independently. With 6 moments and 5 parameters, the system is overidentified with 1 degree of freedom; I report the J-statistic for the overidentifying restriction. Inference is based on a cluster bootstrap at the firm level. I impose the constraints $\varphi_0 \geq 0$, $\varphi_0 + \varphi_1 \geq 0$, and $\max(\varphi_0, \varphi_0 + \varphi_1) < 1/\rho(\tilde{A}L)$ to ensure the scale fixed-point exists.

8.4 Network Regressivity

The structural model formalizes the empirical patterns documented in Sections 6–7 as a theoretical result I call Network Regressivity.

Suppose: (i) $\varphi_S > \varphi_L \geq 0$ (small firms have stronger scale economies); (ii) $\xi_S < 1 < \xi_L$ (input complementarity for small firms, substitutability for large); (iii) large firms have more diversified customer and supplier networks. Then the ratio of small-to-large Domar weight losses is amplified at each stage of the model:

$$\underbrace{\frac{|\Delta \log \lambda_S|}{|\Delta \log \lambda_L|}}_{\text{Domar weight losses}} > \underbrace{\frac{|\theta_S^{eff}|}{|\theta_L^{eff}|}}_{\text{Effective shocks}} > \underbrace{\frac{|\theta_S|}{|\theta_L|}}_{\text{Exogenous shocks}} \tag{22}$$

The first inequality follows from the input substitution channel: with $\xi_S < 1$, small firms cannot substitute toward alternative suppliers, amplifying their Domar weight losses beyond what their effective shocks alone would predict; with $\xi_L > 1$, large firms can reallocate inputs, partially buffering their losses. The second inequality follows from the scale channel: small firms' TFP is more sensitive to customer health ($\varphi_S > \varphi_L$), and their less diversified customer base means higher average customer exposure G_i per unit shock, so their effective shocks are amplified more relative to their exogenous shocks. Together, these two channels create a systematic pattern in which regulatory shocks concentrate losses on small firms—a distributional consequence invisible to models with homogeneous elasticities.

This result also reveals a dual role of firm size: small firms suffer disproportionate losses when exposed to shocks, but because they are peripheral (low Domar weights, few connections), their exit generates small aggregate spillovers. Large firms buffer their own losses through diversification but would generate large cascades if forced to contract. This asymmetry is central to the welfare comparison of policy instruments below.

9 Welfare Analysis

9.1 Social Welfare Function

The welfare analysis evaluates environmental regulation through a social welfare function that balances economic output against pollution damage:

$$W(\boldsymbol{\theta}) = \log \mathrm{GDP}(\boldsymbol{\theta}) - \tau^* E(\boldsymbol{\theta}) \tag{23}$$

where τ^* is the constant marginal social cost of pollution (yuan/ton) and $E(\boldsymbol{\theta}) = \sum_i e_i \cdot y_i(\boldsymbol{\theta})$ is total emissions, with e_i denoting firm i's emission intensity.

Both the marginal cost and marginal benefit of regulation include network effects. The social marginal cost of reducing firm i's output is:

$$\mathrm{MC}_i = \underbrace{\lambda_i}_{\text{Direct (Hulten)}} + \underbrace{\sum_k \varphi_k \lambda_k [\tilde{A}L]_{ki}}_{\text{Scale amplification}} + O(\theta) \tag{24}$$

The first term is the standard Hulten contribution; the second term captures the additional GDP loss from customer-side TFP declines propagated through the network. The social marginal benefit includes both direct abatement and incidental emission reductions from supply chain contractions:

$$\mathrm{MSB}_i = \tau^* \left(\underbrace{e_i \cdot y_i}_{\text{Direct abatement}} + \underbrace{\sum_{j \neq i} e_j \frac{\partial y_j}{\partial \theta_i}}_{\text{Network emission reductions}} \right) \tag{25}$$

Overregulation from ignoring network costs. A key implication is that a regulator who ignores network externalities will overregulate. Environmental regulation targets dirty firms, but their upstream suppliers and downstream customers are systematically cleaner ($e_j \ll e_i$). When a regulated firm contracts, the GDP cascade propagates through the supply chain at full economic weight—each affected firm's Domar weight enters the MC regardless of its emission profile. But the incidental emission reductions are attenuated by the trading partners' lower emission intensities: a logistics provider or equipment manufacturer that loses a client in the regulated zone suffers real GDP losses but generates little emission reduction. Network effects therefore steepen the MC curve more than the MSB curve. A partial-equilibrium regulator who equates the *private* marginal cost (λ_i only) to the *direct* marginal benefit ($\tau^* e_i y_i$ only) will set $|\bar{\theta}|$ too large, overshooting the social optimum (Figure 11).

9.2 Calibrating Pollution Damages

I anchor τ^* on the revealed-preference willingness-to-pay estimates from Ito and Zhang (2020), who exploit residential sorting and housing markets in China to estimate households' marginal WTP for reductions in PM_{10} exposure. Their central estimate implies a social cost of approximately 23,000–47,000 yuan per ton of PM_{10}-equivalent emissions. For robustness, I consider three tiers: a low tier (1,260–6,300 yuan/ton) based on China's Environmental Protection Tax (EPT) rates, a medium tier (23,000–47,000 yuan/ton) based on revealed-preference estimates, and a high tier (50,000–230,000 yuan/ton) based on engineering health damage assessments.

9.3 Imputing Firm-Level Emissions

Since VAT transaction data contain sales but not emissions, I impute firm-level emissions using the Environmental Survey and Reporting Database (ESR), which covers approximately 85% of national industrial emissions and reports firm-level pollutant outputs (SO_2, smoke/dust, COD, ammonia nitrogen) alongside industrial output at the 4-digit industry level.

I proceed in three steps. First, I compute industry-level emission intensity from the 2013 ESR: $e_{s,p} = \sum_{i\in s} E_{ip}^{\text{ESR}} / \sum_{i\in s} Y_i^{\text{ESR}}$ for each industry s and pollutant p. Second, I impute pre-shock emissions for each VAT firm: $\hat{E}_{i,p} = e_{s(i),p} \times y_{i,\text{pre}}$, where $y_{i,\text{pre}}$ is firm i's 2016 sales. Third, I construct a damage-weighted composite: $\hat{E}_i = \sum_p (\tau_p^*/\bar{\tau}) \cdot \hat{E}_{i,p}$, using the ESR smoke/dust variable that maps directly to the PM_{10}-based WTP from Ito and Zhang (2020). The 4-year gap between the 2013 ESR and the 2017 shock introduces classical measurement error, which attenuates estimates conservatively.

9.4 Optimal Regulation Level

Given the overregulation result above, the natural question is: by how much does the optimal regulation level differ from the observed Xiong'an shutdown, and can a different instrument design reduce the welfare cost? The answer depends on the regulator's available instruments. I consider three progressively refined instruments, each yielding a different characterization of the optimum.

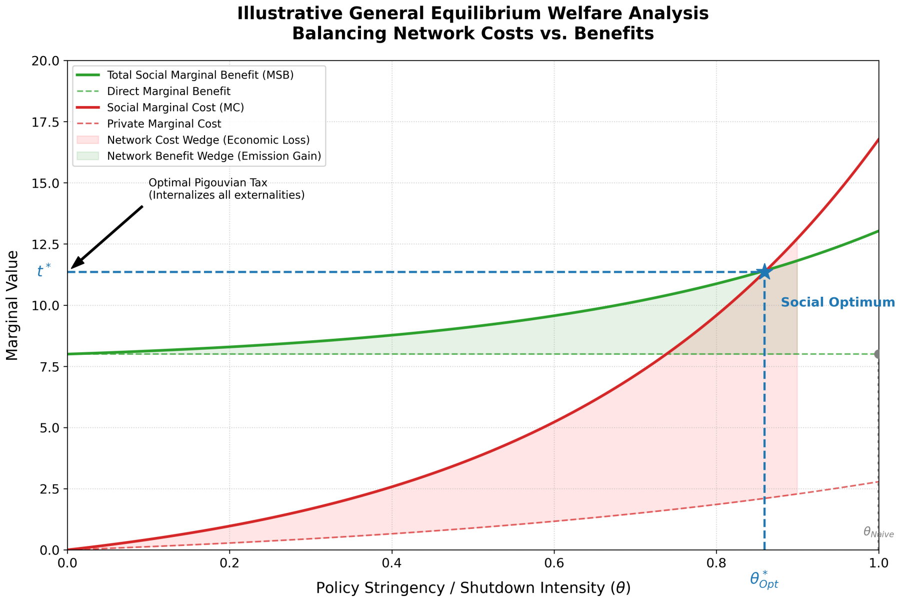

Figure 11: Illustrative Welfare Analysis: Network Costs vs. Benefits

Notes: The figure illustrates the determination of the optimal policy stringency θ^*. The solid red line represents the Social Marginal Cost of abatement (including network losses). The solid green line represents the Total Social Marginal Benefit (including network emission reductions). The optimal policy θ^*_{Opt} (blue star) occurs where full social cost equals full social benefit. The optimal tax t^* aligns private incentives with this social optimum.

Uniform Cap (Quota). The simplest instrument imposes a common proportional output reduction $\theta_i = \bar{\theta}$ on all targeted firms $i \in \mathcal{T}$. The optimal cap $\bar{\theta}^*$ equates social marginal cost (including network propagation losses) to social marginal benefit (emission reductions). The shadow price of the emission constraint at the optimum, t^{cap}, exceeds the direct marginal damage τ^* through a network emission multiplier:

$$t^{cap} = \tau^* \cdot \underbrace{\frac{(\mathbf{E}^0)^\top L\, L^{SE}\, \mathbf{1}_{\mathcal{T}}}{\bar{E}_{\mathcal{T}}}}_{\mathcal{M} > 1} \tag{26}$$

The multiplier $\mathcal{M}$ reflects that regulating targeted firms also reduces emissions through supply chain contractions (the Leontief channel L) and scale-induced TFP declines (the scale channel L^{SE}). Since the social cost per unit of abatement is amplified by network effects, the optimal regulation level under a quota is less stringent than what a partial-equilibrium calculation would suggest.

Uniform Emission Tax. A per-ton tax t on targeted firms allows heterogeneous firm responses:

$$\theta_i^{tax}(t) = \frac{t \cdot E_i}{\lambda_i Y} \propto \frac{E_i}{\lambda_i} \tag{27}$$

Firms with low λ_i/E_i (dirty, low-value) face large output reductions; GDP-critical firms face small reductions. The optimal tax rate t^{Unif} equates the economy-wide marginal cost of further abatement to τ^*, but because the tax sorts firms endogenously, it achieves any given emission target at lower GDP cost than a quota. The welfare gap between the quota and the tax decomposes into three channels:

$$\underbrace{L^Q - L^{Unif}}_{\text{Welfare gap>0}} = \underbrace{\Delta L^{\text{Selection}}}_{\substack{\text{1st order}\\ \text{(standard)}}} + \underbrace{\Delta L^{\text{Convexity}}}_{\substack{\text{2nd order}\\ \textbf{(network-novel)}}} + \underbrace{\Delta L^{\text{Scale}}}_{\substack{\text{1st order}\\ \textbf{(network-novel)}}} \tag{28}$$

The selection channel is standard: the tax concentrates abatement on peripheral, low-value firms by sorting on λ_i/E_i (Buchanan, 1969). The convexity and scale channels are *network-specific*—they arise from production network externalities that neither the firm nor a uniform regulator internalizes. The convexity channel arises because the GDP loss function is convex in θ_i when $\xi < 1$: input complementarity through IO linkages creates second-order interaction terms that make losses accelerate, so heterogeneous reductions generate lower aggregate losses than a uniform cut. Without production linkages, each firm's GDP contribution is linear in its own θ_i (Hulten) and there is no convexity to exploit—this channel is distinct from the Weitzman (1974) prices-vs-quantities result, which arises from uncertainty about marginal costs. The scale channel arises because the quota forces large central firms to contract, triggering demand-dependent TFP declines in their many suppliers; the tax shields these firms, preserving the network's productive capacity. Both novel channels reflect the same two-sided complementarity that generates Network Regressivity, now entering the welfare analysis as reasons favoring price instruments over quantity instruments.

Firm-Specific Taxes (First-Best). Each firm faces a bespoke per-ton tax that internalizes all firm-specific network externalities:

$$t_i^{FB} = \frac{1}{e_i \cdot y_i}\left(\lambda_i + \sum_k \varphi_k \lambda_k [\tilde{A}L]_{ki} + O(\theta)\right) \tag{29}$$

This instrument requires full knowledge of the network structure and serves as a theoretical upper bound. The welfare ordering is $W^{FB} \geq W^{Unif} \geq W^{Cap}$, with strict inequalities when λ_i/E_i varies across targeted firms.

All three instruments pin down the same object—the socially optimal quantity of abatement—but the price-based instruments achieve it more efficiently by allowing firms with different network positions to respond differently. The counterfactual exercises below quantify these welfare gaps.

9.5 Counterfactual Exercises

I conduct three counterfactual exercises using the estimated parameters $\hat{\boldsymbol{\psi}} = (\hat{\sigma}, \hat{\xi}_0, \hat{\xi}_1, \hat{\varphi}_0, \hat{\varphi}_1)$. Figure 11 illustrates the key tradeoff: the optimal policy stringency θ^* occurs where the social marginal cost of abatement (including network losses) equals the total social marginal benefit (including network emission reductions).

CF1: Quota (Baseline). I replicate the Xiong'an shutdown by setting $\theta_i = \hat{\bar{\theta}}$ for all $i \in \mathcal{T}$, where $\hat{\bar{\theta}} \approx -0.53$ is the average log sales decline estimated for Xiong'an polluting firms in Section 5.1. This value reflects the fact that "permanent shutdown" typically refers to the polluting production line rather than the entire firm; many targeted firms retained partial operations. I report the aggregate GDP loss decomposed into the four channels of Equation (17), the scale amplification factor ($\boldsymbol{\lambda}^\top \boldsymbol{\delta} / \boldsymbol{\lambda}^\top \boldsymbol{\theta}$), and the regressivity ratio (small-to-large Domar weight losses).

CF2: No Scale. I impose the same shock vector but set $\varphi = 0$ for all firms. Comparing $\Delta \log Y^Q$ from CF1 with $\Delta \log Y^{\text{no-scale}}$ from CF2 isolates the contribution of the scale channel: the ratio $\Delta \log Y^Q / \Delta \log Y^{\text{no-scale}}$ measures the degree to which demand-dependent TFP amplifies aggregate losses beyond what a standard production network model would predict.

CF3: Uniform Tax. I solve for the uniform tax rate t^{Unif} that achieves the same total emission reduction as the quota ($\Delta E^{tax} = \Delta E^{quota}$), using the continuous firm responses from Equation (27) subject to a $|\theta_i| \leq 1$ corner constraint. Because the tax induces heterogeneous reductions—concentrating abatement on firms with high E_i/λ_i—it achieves the same environmental target at lower GDP cost. I report the welfare gap $L^Q - L^{Unif}$, decomposed into the selection, convexity, and scale channels of Equation (28).

Together, these exercises quantify the aggregate cost of the Xiong'an shutdown, the share of that cost attributable to demand-dependent TFP, and the welfare gain from replacing the observed quantity control with a price mechanism.

10 Conclusion

This paper studies how environmental regulation propagates through production networks and who bears the cost. Exploiting the 2017 Xiong'an pollution shutdown campaign and administrative VAT data covering the universe of firm-to-firm transactions in China, I document three sets of findings.

First, regulatory shocks propagate in both directions through the supply chain. Upstream suppliers at mean exposure experienced an 11.8% decline in sales; downstream customers saw purchases fall by 9.5%. Effects extend to second-degree partners and persist for at least three years. A back-of-envelope aggregation puts first-degree spillover losses at approximately 84% of the direct loss on targeted firms.

Second, the burden of these spillovers falls almost entirely on small firms. Large firms are completely buffered. The mechanism is two-sided complementarity: for small firms, trading partners are gross complements on both the input and output sides, so that losing one partner forces contraction with all others. For large firms, partners are gross substitutes. This asymmetry can be partially traced to the availability of pre-existing alternative partners, which is strongly correlated with firm size. New connection formation is too slow to offset the lost trade, confirming that propagation flows through existing networks.

Third, a general equilibrium model with demand-dependent TFP and size-varying elasticities of substitution formalizes these patterns and draws out their welfare implications. Network externalities create a wedge between the private and social costs of regulation: the social marginal cost of reducing a firm's output includes not only its direct contribution to GDP but also propagated losses through the scale and complementarity channels. This wedge affects optimal regulatory stringency and generates an additional advantage for emissions taxes over quantity controls, distinct from the standard selection-on-abatement-costs mechanism. Under two-sided complementarity, the GDP loss function is convex in the firm-level shock, so that heterogeneous tax-induced reductions achieve the same emission target at lower aggregate cost than a uniform shutdown.

The findings carry implications beyond this particular setting. Any regulation that contracts a firm's operations—whether through taxes, command-and-control mandates, or technology standards—creates upstream demand losses and downstream supply disruptions through the same network channels. The analysis suggests that conventional cost-benefit calculations that ignore network spillovers may substantially understate the true economic burden of environmental regulation, and that the choice between price-based and quantity-based instruments has distributional consequences that extend well beyond the directly regulated firms.

Several directions remain open. The current framework takes the network as given; incorporating endogenous link formation would allow the model to speak to long-run adjustment. The welfare analysis abstracts from dynamic considerations such as investment responses and entry/exit, which may either amplify or attenuate the static losses quantified here. Finally, while the empirical setting provides clean identification, the generalizability of the estimated elasticities—particularly the degree of output side complementarity—is an open question that would benefit from estimation in other settings and shock environments.

References

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012): "The Network Origins of Aggregate Fluctuations," *Econometrica*, 80, 1977–2016.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2020): "The Fall of the Labor Share and the Rise of Superstar Firms," *The Quarterly Journal of Economics*, 135, 645–709.

Baqaee, D. R. and E. Farhi (2019): "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," *Econometrica*, 87, 1155–1203.

——— (2022): "Supply and Demand in Disaggregated Keynesian Economies with an Application to the COVID-19 Crisis," *American Economic Review*, 112, 1397–1436.

Barattieri, A., M. Cacciatore, and N. Traum (2023): "Estimating the effects of government spending through the production network," Tech. rep., National Bureau of Economic Research.

Barrot, J.-N., B. Grassi, and J. Sauvagnat (2021): "Sectoral effects of social distancing," in *AEA Papers and Proceedings*, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, vol. 111, 277–281.

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

Bartram, S. M., K. Hou, and S. Kim (2022): "Real Effects of Climate Policy: Financial Constraints and Spillovers," *Journal of Financial Economics*, 143, 668–696.

Bernard, A. B., E. Dhyne, G. Magerman, K. Manova, and A. Moxnes (2022): "The Origins of Firm Heterogeneity: A Production Network Approach," *Journal of Political Economy*, 130, 1765–1804.

Bernard, A. B., A. Moxnes, and Y. U. Saito (2019): "Production Networks, Geography, and Firm Performance," *Journal of Political Economy*, 127, 639–688.

Bessen, J. (2020): "Industry concentration and information technology," *The Journal of Law and Economics*, 63, 531–555.

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

Buchanan, J. M. (1969): "External Diseconomies, Corrective Taxes, and Market Structure," *The American Economic Review*, 59, 174–177.

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

Castro-Vincenzi, J., G. Khanna, N. Morales, and N. Pandalai-Nayar (2024): "Weathering the storm: Supply chains and climate risk," Tech. rep., National Bureau of Economic Research.

Chen, Q., Z. Chen, Z. Liu, J. C. Suárez Serrato, and D. Y. Xu (2025): "Regulating conglomerates: Evidence from an energy conservation program in china," *American Economic Review*, 115, 408–

447.

Choi, J., A. A. Levchenko, D. Ruzic, and Y. Shim (2024): "Superstars Or Supervillains?: Large Firms in the South Korean Growth Miracle," Tech. rep., National Bureau of Economic Research.

Cui, J. and G. Moschini (2020): "Firm internal network, environmental regulation, and plant death," *Journal of Environmental Economics and Management*, 101, 102319.

Curtis, E. M., W. Miao, F. S. Soliman, J. C. S. Serrato, and D. Y. Xu (2025): "The Geographical Leakage of Environmental Regulation: Evidence from the Clean Air Act," Tech. rep., Technical Report, Technical Report.

De Loecker, J., J. Eeckhout, and G. Unger (2020): "The Rise of Market Power and the Macroeconomic Implications," *The Quarterly Journal of Economics*, 135, 561–644.

Elliott, M., B. Golub, and M. V. Leduc (2022): "Supply network formation and fragility," *American Economic Review*, 112, 2701–2747.

Fell, H. and P. Maniloff (2018): "Leakage in regional environmental policy: The case of the regional greenhouse gas initiative," *Journal of Environmental Economics and Management*, 87, 1–23.

Flaaen, A. and J. Pierce (2024): "Disentangling the Effects of the 2018-2019 Tariffs on a Globally Connected US Manufacturing Sector," *Review of Economics and Statistics*, 1–45.

Fontagné, L. and K. Schubert (2023): "The Economics of Border Carbon Adjustment: Rationale and Impacts of Compensating for Carbon at the Border," *Annual Review of Economics*, 15, 389–424.

Fowlie, M. (2010): "Emissions Trading, Electricity Restructuring, and Investment in Pollution Abatement," *American Economic Review*, 100, 837–869.

Fowlie, M., M. Reguant, and S. P. Ryan (2016): "Market-based emissions regulation and industry dynamics," *Journal of Political Economy*, 124, 249–302.

Fowlie, M. L. (2009): "Incomplete environmental regulation, imperfect competition, and emissions leakage," *American Economic Journal: Economic Policy*, 1, 72–112.

Goulder, L. H. and I. W. H. Parry (2008): "Instrument Choice in Environmental Policy," *Review of Environmental Economics and Policy*, 2, 152–174.

Goulder, L. H., I. W. H. Parry, R. C. Williams, and D. Burtraw (1999): "The Cost-Effectiveness of Alternative Instruments for Environmental Protection in a Second-Best Setting," *Journal of Public Economics*, 72, 329–360.

Gray, W. B. (1987): "The Cost of Regulation: OSHA, EPA and the Productivity Slowdown," *The American Economic Review*, 77, 998–1006.

Greenstone, M. (2002): "The Impacts of Environmental Regulations on Industrial Activity: Evidence from the 1970 and 1977 Clean Air Act Amendments and the Census of Manufactures," *Journal of Political Economy*, 110, 1175–1219.

He, G., S. Wang, and B. Zhang (2020): "Watering down environmental regulation in China," *The quarterly journal of economics*, 135, 2135–2185.

Hoel, M. and L. Karp (2002): "Taxes versus Quotas for a Stock Pollutant," *Resource and Energy*

Economics, 24, 367–384.

Hsieh, C.-T. and E. Rossi-Hansberg (2023): "The Industrial Revolution in Services," *Journal of Political Economy Macroeconomics*, 1, 3–42.

Huneeus, F. (2018): "Production network dynamics and the propagation of shocks," *Graduate thesis, Princeton University, Princeton, NJ*, 52.

Huo, Z., A. A. Levchenko, and N. Pandalai-Nayar (2025): "International Comovement in the Global Production Network," *Review of Economic Studies*, 92, 365–403.

Ito, K. and S. Zhang (2020): "Willingness to Pay for Clean Air: Evidence from Air Purifier Markets in China," *Journal of Political Economy*, 128, 1627–1672.

Jaffe, A. B., S. R. Peterson, P. R. Portney, and R. N. Stavins (1995): "Environmental regulation and the competitiveness of US manufacturing: what does the evidence tell us?" *Journal of Economic literature*, 33, 132–163.

Khanna, G., N. Morales, and N. Pandalai-Nayar (2022): "Supply chain resilience: Evidence from Indian firms," Tech. rep., National Bureau of Economic Research.

Kwon, S. Y., Y. Ma, and K. Zimmermann (2024): "100 years of rising corporate concentration," *American Economic Review*, 114, 2111–2140.

Melitz, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71, 1695–1725.

Miyauchi, Y. (2024): "Matching and Agglomeration: Theory and Evidence From Japanese Firm-to-Firm Trade," *Econometrica*, 92, 1869–1905.

Pigou, A. C. (1920): *The Economics of Welfare*, London: Macmillan.

Walker, W. R. (2013): "The Transitional Costs of Sectoral Reallocation: Evidence from the Clean Air Act and the Workforce," *The Quarterly Journal of Economics*, 128, 1787–1835.

Weitzman, M. L. (1974): "Prices vs. Quantities," *The Review of Economic Studies*, 41, 477–491.

Beyond Targeted Firms: Supply Chain Spillovers of Environmental Regulation in China

Appendix

Table of Contents

A Additional Figures 2

B Additional Tables 6

C Network Propagation Losses Aggregation 13

D Structural Model: Derivations and Discussion 15

D.1 Demand–TFP Equivalence and Model Foundations 15
D.2 Sequential Solution: Formal Argument 16
D.3 Derivation of the GDP Decomposition 17
D.4 Moment Construction: Regression Specifications 19
D.5 Neumann Series Convergence 20

A Additional Figures

Figure A.1: Industry Composition of First- and Second-Degree Network Firms

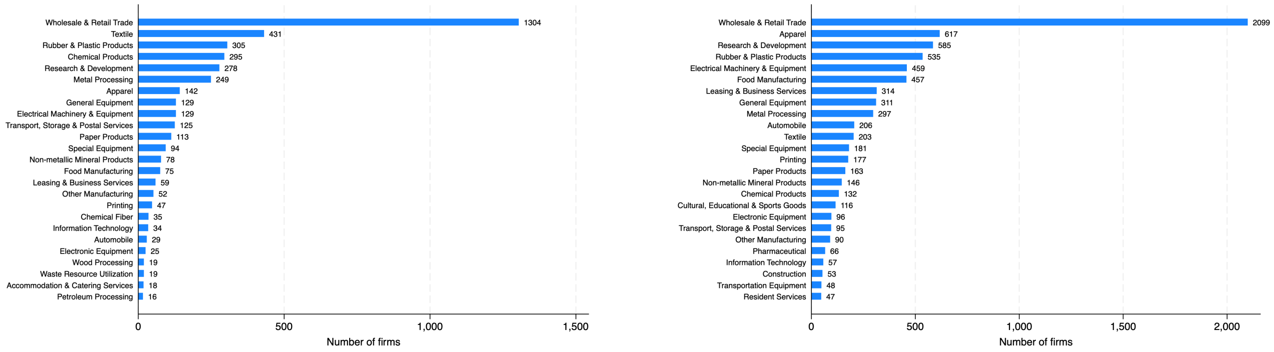

(a) 1st-Degree Upstream Suppliers

(b) 1st-Degree Downstream Customers

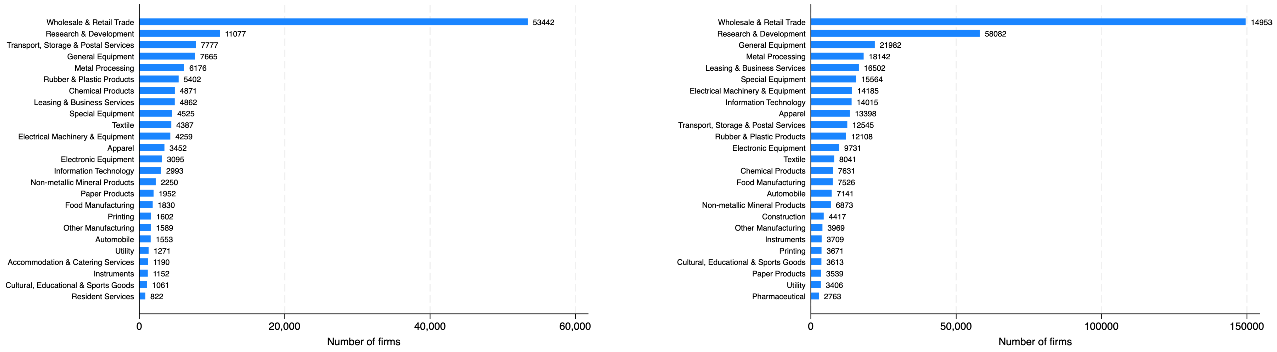

(c) 2nd-Degree Upstream Suppliers

(d) 2nd-Degree Downstream Customers

Notes: This figure plots the number of firms by industry for the first-degree upstream (a), first-degree downstream (b), second-degree upstream (c), and second-degree downstream (d) samples.

Figure A.2: Spatial Distribution of Firms

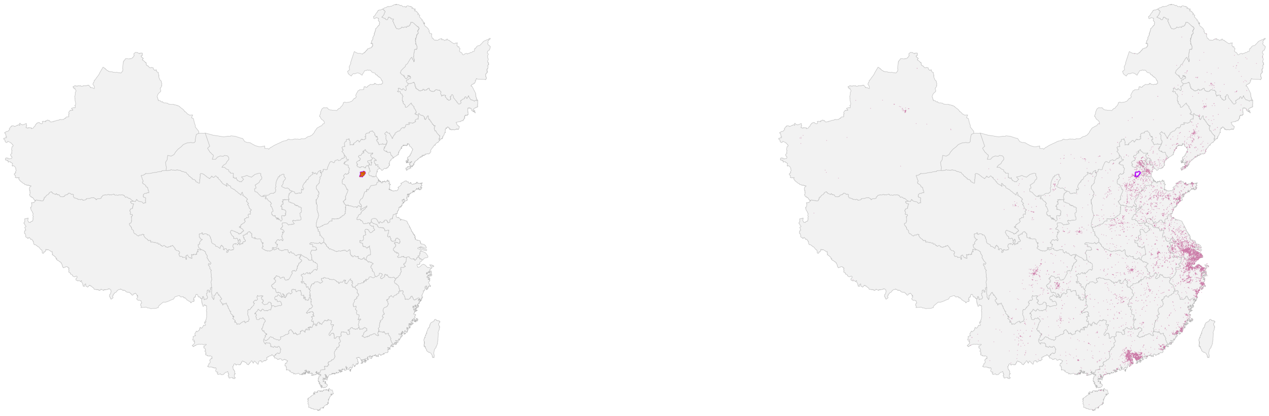

(a) Xiong'an Polluting Firms ($\mathcal{S}$)

(d) Matched Control Firms ($\mathcal{C}$)

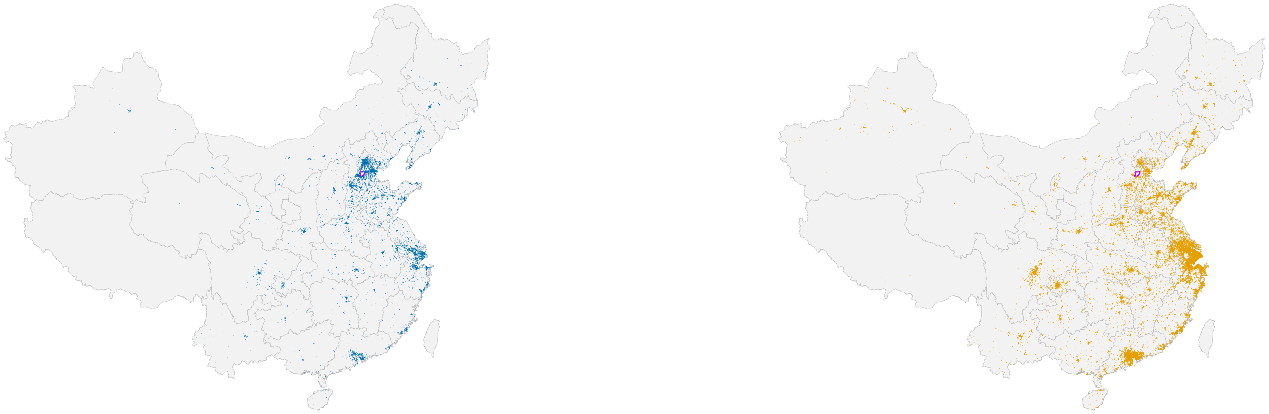

(b) 1st-Degree Partners of $\mathcal{S}$

(e) 1st-Degree Partners of $\mathcal{C}$

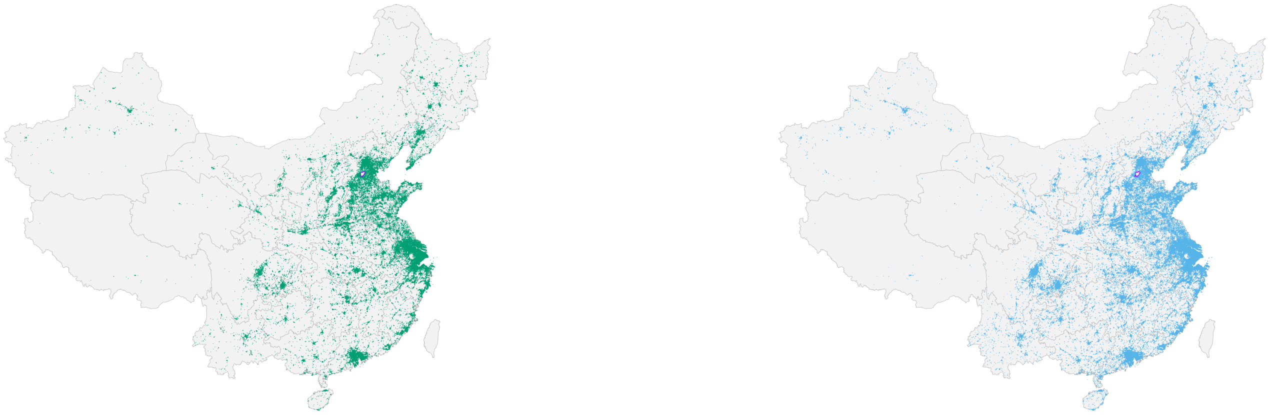

(c) 2nd-Degree Partners of $\mathcal{S}$

(f) 2nd-Degree Partners of $\mathcal{C}$

Notes: This figure maps the geographic location of firms in the analysis samples. The Xiong'an New Area is highlighted with a purple border. Panel (a) shows the Xiong'an polluting firms ($\mathcal{S}$). Panels (b) and (c) show their first- and second-degree trade partners. Panel (d) shows the matched control firms ($\mathcal{C}$). Panels (e) and (f) show the first- and second-degree trade partners of the control firms.

Figure A.3: Distribution of Trade Exposure to Xiong'an Pollution Shutdown Campaign

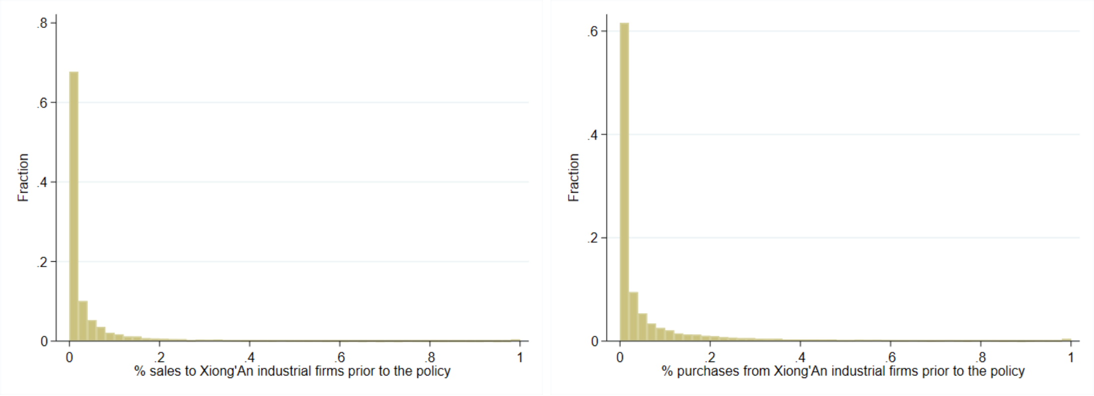

(a) 1st-Degree Upstream Exposure (b) 1st-Degree Downstream Exposure

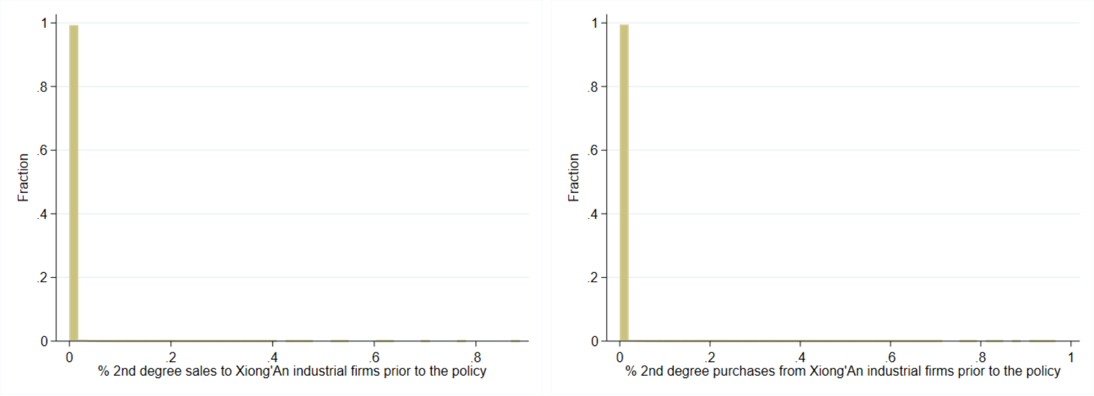

(c) 2nd-Degree Upstream Exposure (d) 2nd-Degree Downstream Exposure

Notes: Panels (a)–(b) plot the first-degree exposure distributions. Panel (a): pre-policy sales share to Xiong'an shutdown firms ($\text{UpExp}^{(1)}$) for the upstream sample. Panel (b): pre-policy purchase share ($\text{DownExp}^{(1)}$) for the downstream sample. Panels (c)–(d) plot the second-degree exposure distributions ($\text{UpExp}^{(2)}$ and $\text{DownExp}^{(2)}$). Trade exposure is computed using the six-month pre-policy window (October 2016–March 2017).

Figure A.4: Event Study Results: New Connection Formation

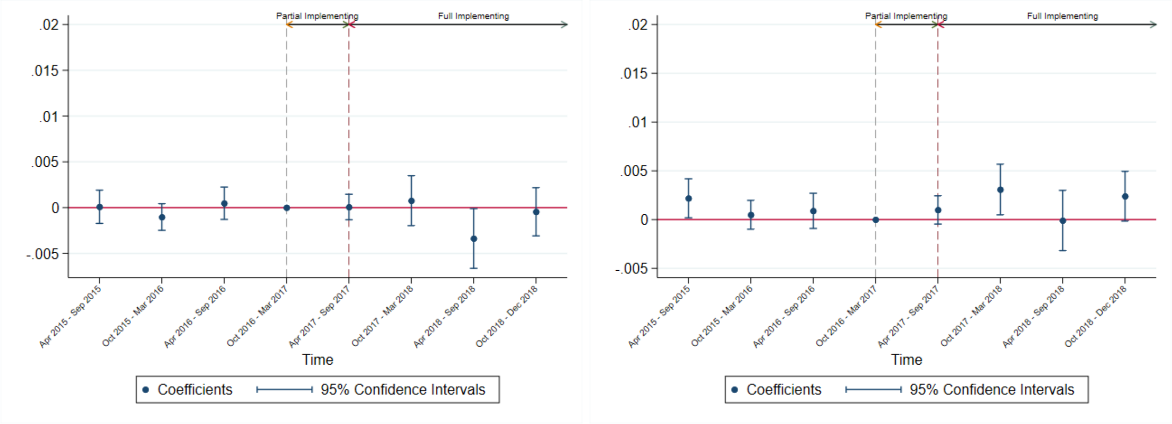

(a) Upstream (Rate of New Customers)

(b) Upstream (Rate of Sales to New Customers)

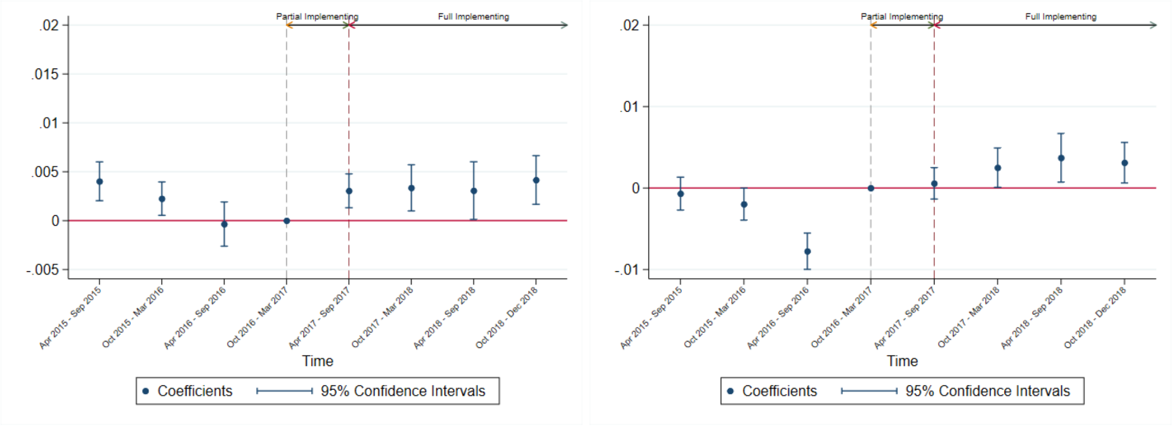

(c) Downstream (Rate of New Suppliers)

(d) Downstream (Rate of Purchases from New Suppliers)

Notes: This figure plots event-study coefficients from the propagation specification (Equation 2) using new connection formation outcomes. New connections are defined as partners that never appeared in any period prior to October 2014–March 2015 (the "legacy" baseline). Denominators are fixed at pre-period levels to avoid endogeneity from contemporaneous partner loss. Pre-period mean outcome values: 0.05 (new customer rate), 0.04 (new sale rate), 0.14 (new supplier rate), 0.07 (new purchase rate). The omitted reference period is October 2016–March 2017 ($\tau = -1$). Standard errors are clustered at the firm level.

B Additional Tables

Table B.1: Summary Statistics of Pre-Policy Firm Characteristics: 2nd Degree Propagation Effects

	(1)					(2)				
Variables	Mean	SD	P5	P95	N	Mean	SD	P5	P95	N
Panel A: Sample of the 2nd degree upstream propagation	2nd Degree Suppliers of Xiong'an Industrial Firms					2nd Degree Suppliers of the Control Firms				
Purchase (Million RMB)	59.97	576.59	0.06	156.28	331,755	18.41	244.13	0.04	48.64	1,180,724
Sale (Million RMB)	67.15	682.62	0.17	166.73	331,755	19.43	275.49	0.14	49.90	1,180,724
Log Purchase	15.00	2.68	11.01	18.87	331,755	14.22	2.58	10.48	17.70	1,180,724
Log Sale	15.40	2.08	12.06	18.93	331,755	14.71	1.81	11.84	17.73	1,180,724
Age	10.51	6.00	3.00	22.00	331,016	9.81	5.72	3.00	20.00	1,178,870
SOE	0.03	0.18	0.00	0.00	331,755	0.02	0.15	0.00	0.00	1,180,724
Distance to the Xiong'an Centrioid (km)	848.79	563.84	103.90	1813.96	331,755	1030.75	524.84	120.43	1827.32	1,180,721
% 2nd Degree Sales to Xiong'an Industrial Firms	0.0007	0.0076	0.0000	0.0020	331,755	-	-	-	-	1,180,724
Panel B: Sample of the 2nd degree downstream propagation	2nd Degree Customers of Xiong'an Industrial Firms					2nd Degree Customers of the Control Firms				
Purchase (Million RMB)	40.37	431.31	0.07	110.81	960,640	8.95	111.93	0.04	24.15	1,105,127
Sale (Million RMB)	43.29	506.27	0.11	111.03	960,640	9.48	120.74	0.07	25.41	1,105,127
Log Purchase	14.94	2.23	11.18	18.52	960,640	14.01	1.96	10.64	17.00	1,105,127
Log Sale	15.01	2.60	11.57	18.53	960,640	14.10	2.49	11.10	17.05	1,105,127
Age	10.10	5.87	3.00	21.00	958,912	9.48	5.61	3.00	20.00	1,103,480
SOE	0.04	0.19	0.00	0.00	960,640	0.02	0.14	0.00	0.00	1,105,127
Distance to the Xiong'an Centrioid (km)	834.45	543.96	105.61	1814.30	960,637	1109.17	524.13	148.21	1832.06	1,105,125
% 2nd Degree Purchases from Xiong'an Industrial Firms	0.0006	0.0078	0.0000	0.0010	960,640	-	-	-	-	1,105,127

Notes: This table presents summary statistics for the pre-policy period (October 2016–March 2017) for the second-degree propagation samples. Panel A shows the second-degree upstream sample (suppliers of first-degree suppliers). Panel B shows the second-degree downstream sample (customers of first-degree customers). "Log Purchase" and "Log Sale" are measured in the pre-policy period. "% 2nd-Degree Sales Exposure" (Panel A) and "% 2nd-Degree Purchase Exposure" (Panel B) are the second-degree exposure measures ($\text{UpExp}^{(2)}$ and $\text{DownExp}^{(2)}$). Distance is measured from each firm to the Xiong'an centroid in kilometers.

Table B.2: Static DID Results: Effects on Annual Log Revenue

	log_Revenue				
	(1)	(2)	(3)	(4)	(5)
XiongAn × Partial	-0.0094				
	(0.1804)				
XiongAn × Full	-0.7450***				
	(0.2444)				
$\text{UpExp}^{(k)}$ × Partial		-0.0037		-0.0011**	
		(0.0118)		(0.0004)	
$\text{UpExp}^{(k)}$ × Full		-0.0496***		-0.0036***	
		(0.0147)		(0.0009)	
$\text{DownExp}^{(k)}$ × Partial			-0.0210***		-0.0001
			(0.0071)		(0.0002)
$\text{DownExp}^{(k)}$ × Full			-0.0426***		-0.0006**
			(0.0097)		(0.0003)
Observations	62,714	1,042,566	1,291,731	10,828,049	16,390,162
R^2	0.6690	0.6100	0.6198	0.5505	0.5332
Firm FE	X	X	X	X	X
Time-Province-Industry FE	X	X	X	X	X
Time-Group FE	X				
Pre-Policy Controls × Time FE		X	X	X	X
Sample	Direct effect	1st Up	1st Down	2nd Up	2nd Down
Mean Exposure		0.04	0.06	0.0007	0.0006

Notes: This table presents the static DID results on annual log revenue. Column (1) shows the direct effect (Eq. 1). Columns (2)–(5) show the propagation effects (Eq. 2). $\text{UpExp}^{(k)}$ and $\text{DownExp}^{(k)}$ are the mean-normalized pre-policy exposure shares for degree $k = 1$ or $k = 2$. “Partial” refers to 2017 and “Full” refers to 2018–2020. Column (1) includes Year-Group fixed effects. Columns (2)–(5) include pre-policy firm characteristics (size, age, ownership) interacted with time fixed effects. Standard errors clustered at the firm level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.3: Robustness Check: Excluding/Duplicating Inter-Group Firms

	Sample: upstream firms				Sample: downstream firms			
	Excluding inter-group firms		Duplicating inter-group firms		Excluding inter-group firms		Duplicating inter-group firms	
	log_purchase	log_sale	log_purchase	log_sale	log_purchase	log_sale	log_purchase	log_sale
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
UpExp$^{(1)}$ × Partial	-0.0104	-0.0499***	-0.0070	-0.0431***				
	(0.0123)	(0.0140)	(0.0112)	(0.0128)				
UpExp$^{(1)}$ × Full	-0.0805***	-0.1251***	-0.0764***	-0.1179***				
	(0.0157)	(0.0206)	(0.0147)	(0.0191)				
DownExp$^{(1)}$ × Partial					-0.0105	-0.0179	-0.0109	-0.0169*
					(0.0094)	(0.0110)	(0.0088)	(0.0102)
DownExp$^{(1)}$ × Full					-0.0964***	-0.0836***	-0.0943***	-0.0798***
					(0.0139)	(0.0153)	(0.0130)	(0.0142)
Observations	1,484,410	1,484,410	1,548,710	1,548,710	1,751,000	1,751,000	1,820,140	1,820,140
R^2	0.7963	0.7243	0.8003	0.7314	0.8052	0.6617	0.8093	0.6660
Firm FE	X	X	X	X	X	X	X	X
Time-Province-Industry FE	X	X	X	X	X	X	X	X
Pre-Policy Controls × Time FE	X	X	X	X	X	X	X	X
Mean 1st degree % sales to XiongAn	0.05	0.05	0.05	0.05				
Mean 1st degree % purchases from XiongAn					0.07	0.07	0.07	0.07

Notes: This table presents robustness checks for the first-degree propagation effects (Table 4) by altering the sample definition for "inter-group" firms. Columns (1)–(2) and (5)–(6) exclude these firms. Columns (3)–(4) and (7)–(8) duplicate these firms. All specifications include firm fixed effects, industry-province-time fixed effects, and pre-policy firm characteristics interacted with time fixed effects. Standard errors clustered at the firm level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.4: Heterogeneous Propagation Effects by Firm Sector

	Sample: upstream firms		Sample: downstream firms	
	log_purchase	log_sale	log_purchase	log_sale
	(1)	(2)	(3)	(4)
$\text{UpExp}^{(1)} \times \text{Full}$	-0.0688***	-0.1236***		
	(0.0202)	(0.0269)		
$\text{UpExp}^{(1)} \times \text{Full} \times \text{Intermediary}$	-0.0162	0.0063		
	(0.0328)	(0.0418)		
$\text{UpExp}^{(1)} \times \text{Full} \times \text{Service}$	-0.0263	0.0430		
	(0.0527)	(0.0735)		
$\text{DownExp}^{(1)} \times \text{Full}$			-0.1093***	-0.0990***
			(0.0190)	(0.0209)
$\text{DownExp}^{(1)} \times \text{Full} \times \text{Intermediary}$			0.0256	0.0313
			(0.0306)	(0.0336)
$\text{DownExp}^{(1)} \times \text{Full} \times \text{Service}$			0.0414	0.0617*
			(0.0340)	(0.0373)
Observations	1,509,850	1,509,850	1,777,280	1,777,280
R^2	0.7973	0.7266	0.8063	0.6622
Firm FE	X	X	X	X
Time-Province-Industry FE	X	X	X	X
Pre-Policy Controls × Time FE	X	X	X	X
Implied effect on trade intermediary	-0.0849***	-0.1173***	-0.0837***	-0.0678**
	(0.0258)	(0.0320)	(0.0239)	(0.0264)
Implied effect on service firms	-0.0951*	-0.0806	-0.0679**	-0.0373
	(0.0486)	(0.0684)	(0.0282)	(0.0308)
Mean 1st degree % sales to XiongAn	0.05	0.05		
Mean 1st degree % purchases from XiongAn			0.07	0.07

Notes: This table presents heterogeneous effects of the first-degree propagation by firm sector, corresponding to Equation 3. The omitted category is the Industrial Sector (Manufacturing, Mining, Construction, Agriculture). “Intermediary” indicates Wholesale/Retail Trade; “Service” indicates all other service sectors. Note that sector-based heterogeneity is partially confounded with firm size, as industrial firms tend to be larger on average. All specifications include firm fixed effects, industry-province-time fixed effects, and pre-policy firm characteristics interacted with time fixed effects. Standard errors clustered at the firm level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.5: Heterogeneous Propagation Effects: Potential Market Partners (7-digit Codes)

	Sample: 1st degree upstream firms		Sample: 1st degree downstream firms	
	log_purchase	log_sale	log_purchase	log_sale
	(1)	(2)	(3)	(4)
$\text{UpExp}^{(1)} \times \text{Full}$	-0.0479**	-0.1073***		
	(0.0221)	(0.0282)		
$\text{UpExp}^{(1)} \times \text{Full} \times S^{EXT}$	-0.0000	0.0000		
	(0.0001)	(0.0001)		
$\text{DownExp}^{(1)} \times \text{Full}$			-0.1032***	-0.0734***
			(0.0192)	(0.0215)
$\text{DownExp}^{(1)} \times \text{Full} \times S^{EXT}$			0.0001	-0.0002
			(0.0002)	(0.0002)
Observations	1,474,430	1,474,430	1,714,400	1,714,400
R^2	0.7965	0.7244	0.8069	0.6617
Firm FE	X	X	X	X
Time-Province-Industry FE	X	X	X	X
Pre-Policy Controls × Time FE	X	X	X	X
Mean 1st degree % sales to XiongAn	0.08	0.08		
Mean 1st degree % purchases from XiongAn			0.12	0.12
Avg. S^{EXT} (Potential Partners, thousand)	307.97	307.97	129.94	129.94

Notes: This table presents heterogeneous effects of the first-degree propagation by the number of potential new partners in the broader market (S^{EXT}), defined using 7-digit product codes. S^{EXT} measures the firm's weighted average number of potential new partners—all firms in the economy that traded the same products but were not pre-existing partners. It is measured in thousands (000s). The interaction coefficient represents the differential effect of having one thousand additional potential partners in the market. The insignificant coefficients (all effectively zero) indicate that market liquidity alone does not help firms replace lost partners, consistent with high frictions in forming new trade relationships and the network stickiness documented in the main text. Average S^{EXT}: 307.97 thousand (upstream), 129.94 thousand (downstream). All specifications include firm fixed effects, industry-province-time fixed effects, and pre-policy firm characteristics interacted with time fixed effects. Standard errors clustered at the firm level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.6: Heterogeneous Propagation Effects: Intensive Margin Substitutability (13-digit Codes)

	Sample: 1st degree upstream firms		Sample: 1st degree downstream firms	
	log_purchase	log_sale	log_purchase	log_sale
	(1)	(2)	(3)	(4)
$\text{UpExp}^{(1)} \times \text{Full}$	-0.0641***	-0.1180***		
	(0.0174)	(0.0221)		
$\text{UpExp}^{(1)} \times \text{Full} \times S^{INT}$	0.0016***	0.0025***		
	(0.0005)	(0.0007)		
$\text{DownExp}^{(1)} \times \text{Full}$			-0.0992***	-0.0872***
			(0.0156)	(0.0173)
$\text{DownExp}^{(1)} \times \text{Full} \times S^{INT}$			0.0183	-0.0016
			(0.0131)	(0.0171)
Observations	1,474,410	1,474,410	1,714,400	1,714,400
R^2	0.7965	0.7244	0.8069	0.6617
Firm FE	X	X	X	X
Time-Province-Industry FE	X	X	X	X
Pre-Policy Controls × Time FE	X	X	X	X
Mean 1st degree % sales to XiongAn	0.08	0.08		
Mean 1st degree % purchases from XiongAn			0.12	0.12
Avg. S^{INT} (Existing Partners)	113.83	113.83	0.89	0.89

Notes: This table presents heterogeneous effects of the first-degree propagation by intensive margin substitutability (S^{INT}), defined using highly granular 13-digit product codes. S^{INT} measures the firm's weighted average number of pre-existing alternative partners for the specific products traded with Xiong'an. The interaction coefficient represents the differential effect of having one additional pre-existing partner for the affected product. Standard errors clustered at the firm level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table B.7: Heterogeneous Propagation Effects: Extensive Margin Substitutability (13-digit Codes)

	Sample: 1st degree upstream firms		Sample: 1st degree downstream firms	
	log_purchase	log_sale	log_purchase	log_sale
	(1)	(2)	(3)	(4)
$\text{UpExp}^{(1)} \times \text{Full}$	-0.0538***	-0.1190***		
	(0.0207)	(0.0277)		
$\text{UpExp}^{(1)} \times \text{Full} \times S^{EXT}$	-0.0000	0.0001		
	(0.0001)	(0.0002)		
$\text{DownExp}^{(1)} \times \text{Full}$			-0.0883***	-0.0666***
			(0.0201)	(0.0222)
$\text{DownExp}^{(1)} \times \text{Full} \times S^{EXT}$			-0.0004	-0.0011
			(0.0008)	(0.0008)
Observations	1,474,410	1,474,410	1,714,400	1,714,400
R^2	0.7965	0.7244	0.8069	0.6617
Firm FE	X	X	X	X
Time-Province-Industry FE	X	X	X	X
Pre-Policy Controls × Time FE	X	X	X	X
Mean 1st degree % sales to XiongAn	0.08	0.08		
Mean 1st degree % purchases from XiongAn			0.12	0.12
Avg. S^{EXT} (Potential Partners, thousand)	121.02	121.02	27.03	27.03

Notes: This table presents heterogeneous effects of the first-degree propagation by extensive margin substitutability (S^{EXT}), defined using 13-digit product codes. S^{EXT} is measured in thousands (000s). The interaction coefficient represents the differential effect of having one thousand additional potential partners in the market for the affected product. Standard errors clustered at the firm level are in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

C Network Propagation Losses Aggregation

This appendix details the methodology for computing aggregate spillover losses from the reduced-form estimates.

Setup. From the size-interacted propagation regression (Equation 3), I estimate:

$$\Delta \ln Y_i = \beta_0 \cdot \text{Exp}_i^{(1)} \times \mathbf{1}\{\text{Full}\} + \beta_1 \cdot \text{Exp}_i^{(1)} \times \mathbf{1}\{\text{Full}\} \times \text{Large}_i + \text{controls} + \varepsilon_i \tag{C.1}$$

where $\text{Exp}_i^{(1)} = \text{RawShare}_i / \overline{\text{RawShare}}$ is the mean-normalized exposure. The coefficient β_0 captures the effect on small firms evaluated at overall-mean normalized exposure ($\text{Exp}_i^{(1)} = 1$), i.e., at the overall sample-mean raw trade share $\bar{E}$.

Exposure adjustment. Since exposure is mean-normalized over the full sample, $\hat{\beta}_0$ reflects the effect at the overall mean raw exposure $\bar{E}$. But small firms have systematically higher raw trade exposure to Xiong'an than the sample average. For a small firm with raw exposure e_i, the predicted log sales change is $\hat{\beta}_0 \times (e_i/\bar{E})$. Evaluating at the small-firm group mean $\bar{E}^S$, the aggregate loss for group $G \in \{\text{Upstream, Downstream}\}$ is:

$$\text{Loss}_G \approx \left(1 - e^{\hat{\beta}_{0,G} \cdot \bar{E}_G^S / \bar{E}_G}\right) \times N_G^S \times \bar{Y}_G^S \tag{C.2}$$

where superscript S denotes small firms. The ratio $\bar{E}_G^S / \bar{E}_G$ captures the fact that small firms are more exposed than the overall sample average: upstream suppliers' mean exposure is 7.7% versus the overall mean of 4.8% (ratio: 1.62), and downstream customers' mean exposure is 12.0% versus 7.1% (ratio: 1.69). Ignoring this size-varying exposure would understate the per-firm loss for the group that actually bears it by approximately 60–70%.

Data inputs and calculation. Table C.1 reports the step-by-step calculation using the current sample statistics.

Table C.1: Detailed Aggregate Loss Calculation

Step	Component	Upstream	Downstream
(a)	$\lvert\hat{\beta}_0^{\text{sale}}\rvert$ (Table 6)	0.1299	0.0926
(b)	Mean raw exposure, small ($\bar{E}^S$)	0.077	0.120
(c)	Overall mean exposure ($\bar{E}$)	0.048	0.071
(d)	Exposure ratio ($\bar{E}^S/\bar{E}$)	1.62	1.69
(e)	Effective log change: (a)×(d)	−0.211	−0.157
(f)	Effective % loss: $1 - e^{(\text{e})}$	19.0%	14.5%
(g)	N^S (small firms)	3,874	7,052
(h)	Mean sales, small ($\bar{Y}^S$, M RMB)	4.88	2.96
(i)	Total small-firm sales (M RMB): (g)×(h)	18,889	20,862
(j)	Total small-firm sales as % of group total	1.2%	1.2%
(k)	Aggregate loss (M RMB): (f)×(i)	3,589	3,024
(l)	**Spillover / Direct loss**	**45.8%**	**38.6%**

Notes: Row (a) is the small-firm propagation coefficient on log sales from the size-interacted specification. Rows (b)–(c) are from summary statistics of the first-degree propagation samples. Row (d) captures the fact that small firms have 62–69% higher trade exposure to Xiong'an than the overall mean. Row (f) uses the exact formula $1 - e^x$ rather than the log-linear approximation. Row (l) uses the direct loss on Xiong'an firms of approximately 7,838M RMB (53.0% of 14,791M total pre-policy sales, where 14,791M = 1,381 firms × 10.71M mean sales from Table 2, Panel A). All sales figures are per semi-annual period in millions of RMB.

Comparison with naive aggregation. For reference, a naive approach that applies the pooled propagation coefficient (Table 4: −11.8% for upstream, −7.9% for downstream) to total group sales would produce a spillover-to-direct ratio exceeding 40×—an implausible magnitude that reflects applying the small-firm elasticity to all firms, assuming uniform exposure, and weighting by the large-firm-dominated sales distribution. Properly accounting for size heterogeneity reduces this ratio to 84%.

D Structural Model: Derivations and Discussion

This appendix collects formal derivations and extended discussion for the structural model in Section 8.

D.1 Demand–TFP Equivalence and Model Foundations

This subsection addresses three related questions about the model's foundations: whether absorbing demand-side feedback into TFP restricts the model, why competitive equilibrium is consistent with CES input demand, and why the demand-dependent TFP formulation is preferred over literal increasing returns to scale.

Demand–TFP equivalence. A concern with the formulation in Section 8.1 is that demand-side feedback is absorbed directly into the TFP term z_i^{eff} rather than being derived from explicit demand functions. Consider two model economies sharing the same primitives (N firms, CES production, competitive markets, Leontief input–output structure) but differing in how the demand–productivity linkage is formalized.

In Model T (the TFP formulation used in the paper), firm i produces $y_i = z_i^{eff} \cdot F(M_i, l_i)$ where z_i^{eff} is taken as given by the firm. In equilibrium, effective TFP satisfies the fixed point $\log z_i^{eff} = \log z_i + \varphi_i \sum_k \tilde{a}_{ik}(L\boldsymbol{\theta}^{eff})_k$.

In Model D (the demand-externality formulation), firm i produces $y_i = z_i \cdot F(M_i, l_i)$ with no demand feedback in TFP. Instead, a demand externality $\mathcal{D}_i$ operates outside the firm's optimization: the firm's measured productivity is $z_i \cdot \mathcal{D}_i(\mathbf{y}_{-i})$ where $\log \mathcal{D}_i = \varphi_i \sum_k \tilde{a}_{ik}(L\boldsymbol{\theta}^{eff})_k$. The firm does not internalize $\mathcal{D}_i$ when choosing inputs.

Since $z_i^{eff} \equiv z_i \cdot \mathcal{D}_i$ by construction, and both models treat the demand-dependent component as a parameter in the firm's cost minimization, the first-order conditions are identical: $c_i^T = c_i^D = c(w, P_i^M)/z_i^{eff}$, and input demands coincide at every price vector. Equilibrium prices satisfy $d \log p_i = -(L\boldsymbol{\theta}^{eff})_i$ in both models, where $\boldsymbol{\theta}^{eff}$ solves the same fixed-point equation. Pre-shock Domar weights coincide because both economies share the same steady state ($\mathcal{D}_i = 1$ before the shock).

By Hulten's theorem, first-order GDP is $d \log Y = \boldsymbol{\lambda}^\top \boldsymbol{\theta}^{eff}$, identical across models. The second-order correction involves changes in Domar weights, $\Delta \log \lambda_i = (\sigma - 1)\Sigma_i(\boldsymbol{\theta}^{eff}) + (\xi_i - 1)\Xi_i(\boldsymbol{\theta}^{eff})$, which depend only on network structure, effective shocks, and substitution parameters—all shared. Hence $d \log Y^T = d \log Y^D$ up to $O(\theta^2)$. $\square$

The equivalence follows the same logic underlying the factor supply ↔ TFP isomorphism in Baqaee and Farhi (2019) and Huo et al. (2025): in efficient economies, the firm's equilibrium conditions depend only on the effective quantity z_i^{eff}, not on the source of the shifter. The equivalence rests on the property that the demand externality is not internalized by the firm. Whether the demand channel is made explicit through $\mathcal{D}_i$ or absorbed into z_i^{eff} does not affect equilibrium outcomes. Under monopolistic competition with variable markups, the equivalence would partially break down: a demand shift and a TFP shift would affect markups through different elasticities. With CES demand (constant markups) the distinction vanishes again; the competitive framework used here avoids the issue entirely.

CES input demand and competitive equilibrium. The CES input demand implied by Equation (6) is smooth and downward-sloping in p_j, which may seem inconsistent with competitive equilibrium. The tension is illusory. Each firm j produces a physically distinct good; the CES aggregator describes a technology—how buyer i combines different inputs—not a demand system over substitute brands. When p_j rises, buyer i substitutes along its isoquant, governed by the technological parameter ξ_i.

The relevant distinction for market structure is whether seller j internalizes its residual demand curve. Under monopolistic competition, firm j would equate marginal revenue to marginal cost and charge a markup $p_j = \frac{\xi}{\xi-1} \cdot MC_j$. Under competitive equilibrium, firm j takes p_j as given and produces at $p_j = MC_j$ with zero profits. Price-taking is sustained because CRS technology implies $AC = MC$ at all output levels; the downward slope in input demand reflects the buyer's substitution technology, not the seller's market power.

Why not literal increasing returns to scale. An alternative formulation would model the output-side feedback as literal IRS: $y_i = z_i \cdot F(M_i, l_i)^{1+\varphi}$ with $\varphi > 0$. This creates three difficulties. First, with degree $1+\varphi > 1$ in the choice variables, marginal cost is decreasing in output and no competitive equilibrium exists; one must instead adopt monopolistic competition. Second, generic IRS does not produce the customer-side feedback documented in Section 7: literal IRS generates scale effects on the input side but has no reason to create feedback from downstream demand conditions. Third, identification breaks down: under literal IRS, the scale parameter governs the curvature of F rather than entering as a multiplicative shifter, and cannot be isolated from σ and ξ using the same ratio-based moments. The demand-dependent TFP formulation avoids all three issues while admitting clean identification from price data. The cost is that the mechanism generating φ is left reduced-form, but as shown above, this is without loss for equilibrium characterization and welfare analysis.

D.2 Sequential Solution: Formal Argument

The GDP decomposition (Equation 17) is computed by first solving the scale fixed-point for $\boldsymbol{\theta}^{eff}$ at pre-shock expenditure shares, then evaluating substitution effects (σ, ξ) conditional on $\boldsymbol{\theta}^{eff}$.

Write GDP as a composite function: $\log Y = f(\boldsymbol{\theta}^{eff}(\boldsymbol{\theta}))$, where f is the standard Baqaee–Farhi GDP function and $\boldsymbol{\theta}^{eff}(\boldsymbol{\theta})$ maps exogenous to effective shocks via the scale fixed point. By the chain rule, the second-order GDP cross-partial decomposes into two sources:

$$\left.\frac{\partial^2 \log Y}{\partial\theta_i\,\partial\theta_j}\right|_{\mathbf{0}} = \underbrace{\sum_{k,l}\left.\frac{\partial^2 f}{\partial\theta_k^{eff}\,\partial\theta_l^{eff}}\right|_{\mathbf{0}}[L^{SE}]_{ki}[L^{SE}]_{lj}}_{\text{Source A}} + \underbrace{\sum_k \lambda_k \left.\frac{\partial^2\theta_k^{eff}}{\partial\theta_i\,\partial\theta_j}\right|_{\mathbf{0}}}_{\text{Source B}} \tag{D.1}$$

Source A captures the standard Baqaee–Farhi reallocation transmitted through L^{SE}—both the sequential and simultaneous procedures agree on this term, which produces channels (c) and (d). Source B captures the curvature of the scale mapping.

Under the sequential procedure, $\boldsymbol{\theta}^{eff} = L^{SE}(\boldsymbol{\omega}^0)\boldsymbol{\theta}$ is linear in $\boldsymbol{\theta}$ (pre-shock shares are fixed), so $\partial^2\theta_k^{eff}/\partial\theta_i\partial\theta_j = 0$ and Source B vanishes. Under the simultaneous solution, expenditure shares respond to prices, making $\boldsymbol{\theta}^{eff}(\boldsymbol{\theta})$ nonlinear. Implicit differentiation of the fixed point shows that Source B scales as $\varphi\cdot\varepsilon\cdot\theta^2$ where $\varepsilon = \max\{|\sigma-1|, |\xi-1|\}$. This is formally $O(\theta^2)$—the same order as the kept second-order terms—but attenuated by φ. For empirically plausible parameters ($\varphi \approx 0.2$, $\varepsilon \approx 0.3$, $\theta \approx 0.1$), Source B is approximately 0.6% of total GDP loss. The error vanishes identically when $\varphi = 0$ or $\sigma = \xi = 1$.

The practical resolution separates estimation from counterfactual computation, following Carvalho et al. (2021). Parameters are estimated using the tractable sequential procedure (Stage 1), since the slope coefficients identifying σ, ξ, and φ come from micro moments (DID regression slopes) that do not depend on the GDP formula. For welfare counterfactuals (Stage 2), the simultaneous equilibrium is solved numerically by iterating: solve $\boldsymbol{\theta}^{eff}$ at current shares, compute prices, update shares via CES formulas, and repeat until shares converge. The converged solution incorporates Source B and all higher-order terms.

D.3 Derivation of the GDP Decomposition

Assumptions. The derivation requires the following:

1. Competitive equilibrium: firms are price-takers with CRS technology.
2. CES production (Equations 5 and 6) with elasticities σ (factors) and ξ_i (inputs).
3. Cobb-Douglas household preferences with fixed expenditure shares $\{\beta_i\}$.
4. The scale fixed-point $(I - \text{diag}(\boldsymbol{\varphi})\tilde{A}L)$ is invertible, guaranteed by $\max_i(\varphi_i) < 1/\rho(\tilde{A}L)$.
5. Shocks $\boldsymbol{\theta}$ are "small" in the sense that the second-order Taylor approximation is adequate.

First-order terms. By Hulten's theorem, GDP responds to effective TFP shocks at first order:

$$d\log Y = \sum_i \lambda_i\,\theta_i^{eff} + O(\theta^2) = \boldsymbol{\lambda}^\top\boldsymbol{\theta}^{eff} \tag{D.2}$$

Substituting $\boldsymbol{\theta}^{eff} = \boldsymbol{\theta} + \boldsymbol{\delta}$ where $\boldsymbol{\delta} = (L^{SE} - I)\boldsymbol{\theta}$ separates the direct Hulten term from the scale amplification:

$$\boldsymbol{\lambda}^\top \boldsymbol{\theta}^{eff} = \underbrace{\boldsymbol{\lambda}^\top \boldsymbol{\theta}}_{\text{(a) Direct}} + \underbrace{\boldsymbol{\lambda}^\top \boldsymbol{\delta}}_{\text{(b) Scale}} \tag{D.3}$$

Second-order terms. At second order, Domar weights shift in response to relative price changes. From the CES cost function, log-differentiating the intermediate share of firm j:

$$d\log \mu_j = (1-\sigma)(1-\mu_j)\left(d\log P_j^M - d\log w\right) \tag{D.4}$$

$$d\log \omega_{jk} = (1-\xi_j)\left(d\log p_k - d\log P_j^M\right) \tag{D.5}$$

where ω_{jk} is firm j's expenditure share on input k, and P_j^M is firm j's intermediate price index. The change in the Domar weight of firm j is:

$$\Delta\lambda_j = \lambda_j \left[(\sigma - 1)\Sigma_j(\boldsymbol{\theta}^{eff}) + (\xi_j - 1)\Xi_j(\boldsymbol{\theta}^{eff})\right] \tag{D.6}$$

where the network statistics are defined in the main text (Section 8.2):

$$\Sigma_j = \mu_j(1-\mu_j)\left(d\log P_j^M - d\log w\right)^2 \tag{D.7}$$

$$\Xi_j = \sum_k \omega_{jk}\left(d\log p_k - d\log P_j^M\right)^2 \tag{D.8}$$

The price changes entering these expressions are $d\log \mathbf{p} = -L\boldsymbol{\theta}^{eff}$, so Σ_j and Ξ_j are determined by $\boldsymbol{\theta}^{eff}$ and the pre-shock network structure.

The second-order GDP correction consists of two parts. Channel (c), the Domar reallocation:

$$\frac{1}{2}\sum_j \Delta\lambda_j \cdot \theta_j^{eff} = \frac{1}{2}\sum_j (\lambda_j^* - \lambda_j^0)\theta_j^{eff} \tag{D.9}$$

Channel (d), the intermediate share adjustment, captures the additional second-order effect from changes in the economy-wide intermediate share when $\sigma \neq 1$:

$$\frac{1}{2}\mu(\sigma - 1)(\boldsymbol{\theta}^{eff})^\top \boldsymbol{\Lambda}(I - A)L\,\boldsymbol{\theta}^{eff} \tag{D.10}$$

following the derivation in Carvalho et al. (2021). Summing channels (a)–(d) yields Equation (17).

D.4 Moment Construction: Regression Specifications

This subsection provides the full panel regression specifications used to construct the empirical moments for the minimum-distance estimator.

Moments 1–2 (Identifying σ). The structural relationship (Equation 18) is estimated in a panel DID where the dependent variable is the log purchase-to-sales ratio and the regressor is the change in the firm's intermediate price index:

$$\log \frac{\text{Purch}_{it}}{\text{Sales}_{it}} = \delta_i + \gamma_{sct} + \beta_\sigma \left(d \log P_i^M \times \text{Post}_t \right) + \varepsilon_{it} \tag{D.11}$$

where δ_i are firm fixed effects and γ_{sct} are sector×county×time fixed effects. The intermediate price index $d \log P_i^M = \sum_j \omega_{ij,\text{pre}}\, d \log p_j$ is constructed using pre-shock expenditure weights and observed transaction-level prices. Fixed effects absorb the wage change, so $\hat{\beta}_\sigma$ identifies $(1-\sigma)(1-\mu)$. Since σ is common across firm sizes, I pool small and large firms within each position group (no size interaction). I compute two moments: $\hat{m}_1$ (pooled downstream) and $\hat{m}_2$ (pooled upstream), providing one overidentifying restriction.

Moments 3–4 (Identifying ξ). The structural relationship (Equation 19) is estimated as:

$$\log \frac{\text{Purch}_{it}^{AFF}}{\text{Purch}_{it}^{OTH}} = \delta_i + \gamma_{sct} + \beta_\xi \left(\Delta p_i^{GAP} \times \text{Post}_t \right) + \beta_{\xi,L} \left(\Delta p_i^{GAP} \times \text{Large}_i \times \text{Post}_t \right) + \varepsilon_{it} \tag{D.12}$$

where $\Delta p_i^{GAP} \equiv \overline{\Delta \log p_i}^{AFF} - \overline{\Delta \log p_i}^{OTH}$ is the within-firm gap in average supplier price changes between affected and unaffected groups, computed with pre-shock Laspeyres weights. The moments are $\hat{m}_3 = \hat{\beta}_\xi$ (identifying $1 - \xi_0$) and $\hat{m}_4 = \hat{\beta}_\xi + \hat{\beta}_{\xi,L}$ (identifying $1 - \xi_0 - \xi_1$).

Moments 5–6 (Identifying φ). The structural relationship is estimated for upstream firms not themselves targeted ($\theta_i = 0$):

$$\log p_{it} = \delta_i + \gamma_t + \delta_1^p \left(G_i \times \text{Post}_t \right) + \delta_2^p \left(G_i \times \text{Large}_i \times \text{Post}_t \right) + \beta^p \left(d \log P_i^M \times \text{Post}_t \right) + \varepsilon_{it} \tag{D.13}$$

where $G_i = [\tilde{A} L \boldsymbol{\theta}]_i$ is the firm's sales-share-weighted customer exposure and $d \log P_i^M$ controls for input price pass-through. The control for $d \log P_i^M$ is necessary to isolate the scale channel: without it, δ_1^p would confound the TFP erosion (φ) with the mechanical pass-through of higher input costs. The moments are $\hat{m}_5 = \hat{\delta}_1^p$ (identifying $-\varphi_0$) and $\hat{m}_6 = \hat{\delta}_1^p + \hat{\delta}_2^p$ (identifying $-(\varphi_0 + \varphi_1)$).

D.5 Neumann Series Convergence

The scale-extended Leontief inverse $L^{SE} = (I - \text{diag}(\boldsymbol{\varphi})\tilde{A}L)^{-1}$ exists when the spectral radius of $\text{diag}(\boldsymbol{\varphi})\tilde{A}L$ is strictly less than one. A sufficient condition is $\max_i(\varphi_i) < 1/\rho(\tilde{A}L)$, where $\rho(\cdot)$ denotes the spectral radius.

Under this condition, L^{SE} admits the convergent Neumann expansion:

$$L^{SE} = \sum_{n=0}^{\infty} [\text{diag}(\boldsymbol{\varphi})\tilde{A}L]^n \tag{D.14}$$

The n-th term captures n rounds of customer-side feedback: the zeroth-order term is the identity (exogenous shock only), the first-order term captures the direct scale response to customer contraction, and higher-order terms capture cascading amplification through chains of customer relationships. Since $\tilde{A}$ has rows summing to at most one (sales shares) and L has bounded entries (the Leontief inverse of a productive economy), $\rho(\tilde{A}L)$ is finite. In practice, the estimated $\hat{\varphi}$ values are well below this bound, and the series converges within three to five iterations.