Transmission of Financial Shocks through Production Networks: Theory and Empirical Evidence from China's Real Estate Deleveraging*

Haoxuan Su† Anxu Wang‡ Yang Yao§

August 26, 2026

Abstract

How far does financial regulation aimed at one sector travel through the wider economy? We study China's 2020 "Three Red Lines" policy, which tied real estate developers' permitted debt growth to three balance-sheet thresholds, and trace its reach along the trade-credit links of the production network. Combining listed firms' balance sheets with China's input-output table, we document both sides of the trade-credit relationship: payables rose at constrained developers and receivables at their suppliers—evidence that trade credit is the financial friction on the network's own transaction edges rather than a parallel credit channel. We then develop a multi-sector model in which each firm chooses what fraction of its input bill to prepay, subject to a binding borrowing constraint, so payment terms respond endogenously to financial conditions. Calibrated to China, the model generates a GDP trough of -4.7%; roughly 40% of the cumulative output response is attributable to the endogenous adjustment of trade credit, a $1.7\times$ amplification over fixed payment terms. Sector-specific financial regulation should therefore be judged by its network spillovers, not only by its effect on the targeted sector.

Keywords: Trade credit, production networks, financial shocks, real estate, China
JEL Codes: E32, E44, G28, L14, R31

*We thank seminar participants at Peking University and Fudan University, and participants at the China Economics Annual Conference Summer Forum at Dali University, for helpful comments. The authors received no specific funding for this research and have no relevant financial relationships or other potential conflicts of interest to disclose.

†China Center for Economic Research, National School of Development, Peking University (hxsu2023@nsd.pku.edu.cn).

‡China Center for Economic Research, National School of Development, Peking University (axwang2023@nsd.pku.edu.cn).

§Dishui Lake Advanced Financial Institute, Shanghai University of Finance and Economics; China Center for Economic Research, National School of Development, Peking University (yyao@nsd.pku.edu.cn).

1 Introduction

Financial regulation aimed at a single sector has become a standard instrument. Down-payment floors on million-dollar homes in Toronto and Vancouver, loan-to-value limits adopted across dozens of countries, dynamic provisioning for Spanish banks, and China's 2020 "Three Red Lines" rules for property developers all constrain the credit of one set of borrowers and no other (Jiménez et al., 2017; DeFusco et al., 2020; Han et al., 2021; Acharya et al., 2022). Evaluations of such regulation accordingly stop at the market the rule targets. Two findings suggest that the stopping point is premature. Shocks to one sector are amplified economy-wide by linkages through the production network (Long and Plosser, 1983; Acemoglu et al., 2012; Carvalho, 2014; Baqaee and Farhi, 2019; Ozdagli and Weber, 2026) and by financial interconnections (Acemoglu et al., 2015; Huremović et al., 2026). And targeted financial constraints reshape behavior within the very market they regulate: Han et al. (2021) show that a down-payment requirement applied to homes above a price threshold produced bunching of listings and intensified bidding at the threshold rather than a simple contraction in sales. If the constrained parties' response extends beyond their own market—if a substantial share of the aggregate cost travels through payment terms between firms, terms that regulators neither observe nor control—then judging such a policy by its effect on the targeted sector understates its cost, and a macroprudential framework that holds payment terms fixed will miss a large part of the response. How large a part is a quantitative question.

We answer it in arguably the largest sector-targeted deleveraging yet attempted: China's "Three Red Lines" policy, announced in August 2020, which tied property developers' permitted debt growth to three balance-sheet thresholds. When a developer's access to credit is capped, the squeeze does not stop at its construction sites. The developer keeps ordering steel but pays for it later; receivables pile up on the steel supplier's books; the supplier, now short of cash, stretches its own payments to the firms behind it. No loan contract connects these firms—only invoices do. This paper argues that such chains of stretched invoices carried a substantial share of the shock's aggregate cost. In the data, payables surged at constrained developers while receivables surged at their suppliers. In a model calibrated to China's production network, the endogenous response of trade credit carries roughly 40% of the aggregate loss—a 1.7$\times$ amplification relative to an economy with fixed payment terms.

Reaching that answer takes three ingredients. First, a quarterly panel of Chinese listed firms provides balance sheets on both sides of the trade-credit relationship: the payables of developers and the receivables of their suppliers. Second, China's 2020 input-output table supplies the network: we measure each industry's exposure to real estate by the Leontief inverse, which captures direct and higher-order supply links, and we interact that exposure with each firm's predetermined reliance on trade credit. Third, a dynamic multi-sector model in which every firm chooses what fraction of its input bill to prepay—the prepayment fraction ϕ—subject to a binding borrowing constraint turns the reduced-form evidence into an aggregate answer: when the developer's constraint tightens, reduced prepayment passes the financial pressure upstream, and we can decompose the equilibrium response into the part carried by trade credit and the part the network would transmit on its own. This interaction is the empirical counterpart of the model's transmission term: the reduced-form coefficient measures the same equilibrium object the endogenous prepayment fraction generates when a developer's constraint tightens.

Trade credit is large enough to carry such a shock. Deferred payment between trading partners accounts for over half of firms' short-term liabilities across OECD countries; in the United States it exceeds bank loans by a factor of three and commercial paper by a factor of fifteen (Altinoglu, 2021). In China, accounts payable make up roughly one fifth of listed firms' total liabilities, and a larger share for non-state firms (Cun et al., 2022). An extensive empirical literature shows that trade credit responds to conditions in the market for bank credit (Petersen and Rajan, 1997; Cuñat, 2007; Garcia-Appendini and Montoriol-Garriga, 2013) and passes financial distress along supply chains (Boissay and Gropp, 2013; Jacobson and von Schedvin, 2015; Costello, 2020; Almeida et al., 2024). The mechanism is simple: a firm that delays paying its suppliers converts its own liquidity shortage into a cut in the supply of credit to its trading partners, even while goods keep flowing. Two results frame what we should expect. Adelino et al. (2023) study the mirror image of our experiment: when the European Central Bank's bond purchases eased financing for eligible firms, those firms passed the easing downstream by extending more trade credit, and their customers raised investment and employment. Payment terms, that is, carry monetary as well as commercial content, and they carry it along the production network. Murfin and Njoroge (2015) establish the cost side of the same margin from the buyer's end: large buyers that stretch payment terms finance themselves from their smaller suppliers, and the suppliers absorb the cost in lower invest-

ment, the more so the tighter their own financing constraints. Our setting reverses the sign of the shock and the direction of travel—a contraction in bank credit to buyers, transmitted upstream by slower payment—and asks what the resulting chain costs in the aggregate. Direct evidence on this transmission nonetheless remains scarce, because it requires observing payment behavior between specific firm pairs—data that are rarely available, in China or elsewhere.

Most evidence on the transmission of financial shocks comes from aggregate events—crises and monetary tightenings—in which every firm is hit at once, leaving neither a clean control group nor a traceable origin. The Three Red Lines policy provides an unusually clean setting in which to observe the channel. The announcement came with little advance warning, limiting anticipation effects. The policy targeted a well-defined sector—real estate development—creating a sharp boundary between treated and untreated firms, and within the treated sector its tiered thresholds generated dose-response variation across developers. Real estate sits near the center of China's production network, with deep upstream links to construction, basic manufacturing, and services (Rogoff and Yang, 2021), so the spillovers we document reflect a macroeconomically relevant channel rather than a curiosity. And the shock was large enough to measure yet contained enough that transmission can be attributed to identifiable supply-chain links rather than to a generalized financial panic.

Our findings come in four steps. First, we establish the economy-wide reach of the shock. For every listed firm we build an exposure measure that multiplies its pre-policy ratio of receivables to assets—how much credit the firm extends to its customers—by its industry's Leontief linkage to real estate. After the policy, a firm one percentage point more exposed lost 236 million RMB of cash holdings and 204 million RMB of quarterly sales; the losses appear immediately after the announcement, deepen over time, and persist for more than three years. The estimates are robust to—indeed strengthen under—a predetermined size control that interacts pre-policy employment with the post-policy indicator.

Second, we verify that the targeted firms transmitted, rather than absorbed, the squeeze. After the policy, developers' short-term bank loans—the margin of credit the policy directly capped—fell by 678 million RMB and their cash holdings by 500 million RMB relative to other listed firms, while their accounts payable rose by 5.58 billion RMB: as bank credit dried up, developers paid their suppliers later, shifting the liquidity shortage upstream. The regulation's own thresholds

confirm the dose. Within the real estate sector, developers that breached at least one red line on the eve of the policy increased payables by 8.71 billion RMB relative to compliant developers.

Crucially, the channel is visible from both ends of the same transactions. Among listed suppliers that disclose a real estate developer as a major customer, receivables rose by 11.5 percentage points of sales and log sales fell by 0.229; across the full sample of non-real-estate firms, a percentage point of input-output exposure to real estate reduced cash holdings by 70 million RMB and debt by 123 million RMB. The payables piling up at developers and the receivables piling up at their suppliers are two sides of the same trade-credit relationship, measured on different firms through independent identification strategies. A pure contraction in demand would shrink volumes on both sides but need not, by itself, generate the offsetting rise in payables at developers and receivables at their suppliers that we observe. The two-sided pattern thus points to the friction itself.

Trade credit, in short, is the financial friction on the production network's own transaction edges, not a parallel credit channel: the network supplies the edges, and the friction on them is what carries a financial shock.

Finally, we ask what these firm-level responses add up to. We embed the prepayment choice in a seven-sector dynamic model calibrated to China's 2020 input-output table, in which real estate produces a durable asset for households and finances its inputs subject to a sector-specific borrowing limit. A tightening of that limit—peaking at 25.9% and calibrated to developers' observed credit contraction between 2019 and 2022—generates a GDP trough of -4.7% and a cumulative twenty-quarter loss of 31.9 percentage-point-quarters. Holding every firm's prepayment fraction at its steady-state value cuts the cumulative loss to 19.1: roughly 40% of the aggregate response is attributable to the endogenous adjustment of trade credit, an amplification of $1.7\times$. As an over-identifying check, an accounting aggregation of our firm-level sales estimates implies an economy-wide loss of the same order of magnitude as the model's—two largely independent routes—resting on different identifying variation—that nonetheless converge.

Three features of the design address the main identification concerns. Exposure is predetermined: the firm-level component is measured entirely before the announcement and shows no pre-policy trend, and event-study coefficients are flat before the third quarter of 2020. Industry-by-time fixed effects absorb every industry-wide shock—including each industry's direct demand

exposure to the housing downturn—so the estimates compare firms within the same industry that differ in their reliance on trade credit; and the comparison within real estate holds the housing market itself fixed, since breaching and compliant developers faced the same demand conditions and differed only in regulatory bite. And the channel's incidence is structured in a way that is hard to generate mechanically: it runs almost entirely through higher-order network paths rather than direct sales to real estate, and it concentrates among financially constrained private firms, while state-owned firms substitute toward bank borrowing.

This paper makes three contributions. First, we provide two-sided evidence that trade credit transmits a sector-specific financial shock through a production network. Comprehensive data on credit between firm pairs rarely exist; we substitute the combination of firm balance sheets, disclosed customer relationships, and input-output structure, and we observe the channel at both ends—payment delays at the constrained buyers and receivable accumulation at their suppliers. Among concurrent studies of the Three Red Lines, the closest is Chen et al. (2024), who measure exposure through stock-return correlations and document spillovers in stock prices, bond spreads, and investment; our exposure measure captures commercial rather than financial connections to real estate, and the two-sided demand–supply evidence is unavailable from a single-sided design. Ma and Wang (2025) also studies the policy but focuses on developers' own financial risk rather than on transmission beyond the sector.

Second, we develop a model of endogenous, constraint-driven trade credit. Its central object is the buyer's prepayment choice: each firm decides what fraction of its input bill to pay in cash and what fraction to defer, and that fraction is pinned down by the shadow value of the firm's binding borrowing constraint. Payment terms thereby become equilibrium objects that move with financial conditions—which is exactly why they carry a financial shock, rather than a productivity shock, upstream. It is also where our evidence on heterogeneity points: the spillover concentrates in financially constrained firms, the firms for which the shadow value moves most. And because it grounds payment terms in a binding constraint and solves the network transparently, the framework applies wherever a sector-specific financial shock meets a production network—regulatory deleveraging elsewhere, disruptions to bank credit (Huremović et al., 2026), or distress at a firm near the center of a supply network.

Third, our quantification of how far the Three Red Lines shock traveled carries a lesson for

regulatory design: a leverage cap aimed at one sector propagates economy-wide along the trade-credit edges of the production network, so such regulation must be judged by its network spillovers, not merely by its effect on the targeted firms. For a sector as deeply embedded as Chinese real estate, a large share of the aggregate cost arises outside it.

Our model is closest to Altinoglu (2021), who first brought trade credit into a quantitative model of the production network, modeling it as supplier credit collateralized by downstream revenue. Because Altinoglu (2021) treats the trade-credit share as exogenous, endogenizing the prepayment fraction lets us (i) generate the co-movement of payables and receivables as an equilibrium response to a financial shock, (ii) decompose the contribution of the friction from pure Leontief propagation, and (iii) quantify the spillover of a leverage cap aimed at one sector—the policy experiment that motivates the paper. A complementary line derives trade-credit terms from a choice between bank and supplier credit (Reischer, 2020; Luo, 2020; Cun et al., 2022). Both approaches are largely static, and in the cost-of-credit line the operative margin is a price: payment terms move when interest rates or collateral values move. The Three Red Lines policy instead capped quantities, and it did so in a dynamic economy; the buyer-side margin we model, in which a binding borrowing limit moves the prepayment fraction period by period, is the margin such a shock activates.

The paper proceeds as follows. Section 2 provides institutional background on the Three Red Lines policy and on real estate's position in China's economy. Section 3 presents the reduced-form evidence. Section 4 formalizes the mechanism—first in a simple two-sector model that states it as a proposition, then in a multi-sector network calibrated to China's input-output table—and reports the baseline impulse response. Section 5 presents the quantitative analyses: the decomposition of the trade-credit channel, the over-identification check against the empirical estimates, the sectoral decomposition, and policy counterfactuals. Section 6 concludes.

2 Institutional Background

This section sets out the institutional context for our analysis in three steps: the systemic position of real estate in China's economy, the design and timing of the Three Red Lines policy, and the financing practices—above all the reliance of developers on trade credit—that primed the sector's

supply chain to transmit the shock.

2.1 The Real Estate Sector in China's Economy

Real estate occupies a central position in China's economy. While the sector's direct value-added—comprising real estate services and construction—accounts for roughly 12% of GDP, its broader economic footprint is substantially larger. When upstream intermediate inputs are included, real estate and related activities contribute an estimated 25–29% of GDP (Rogoff and Yang, 2021). The sector is a major consumer of steel, cement, glass, and other building materials, and its demand ripples through machinery, chemicals, furniture, and home appliances.

Beyond production linkages, real estate is deeply embedded in China's fiscal and financial architecture. Sales of land-use rights have historically constituted over a third of local government revenue at the height of the property cycle, making the sector a fiscal pillar for subnational governments. On the household side, housing represents approximately 60–70% of urban household assets (People's Bank of China, 2020), far exceeding the corresponding figure in most developed economies. Because household wealth is so concentrated in real estate, developments in the housing market have immediate implications for consumption, saving behavior, and financial stability.

The sector's systemic importance implies that shocks originating in real estate can propagate broadly through the economy. Yet prior to the events we study, the empirical magnitude of such propagation—and the channels through which it occurs—remained poorly understood.

2.2 The Three Red Lines Policy

In August 2020, the Ministry of Housing and Urban-Rural Development (MOHURD) and the People's Bank of China (PBOC) convened a meeting with twelve major real estate developers to announce a new regulatory framework for managing developer leverage. The policy, known as the "Three Red Lines", imposed three financial thresholds on real estate developers:

1. Liability-to-asset ratio (excluding advance receipts from pre-sales) must not exceed 70%;
2. Net debt-to-equity ratio must not exceed 100%;
3. Cash-to-short-term-debt ratio must be at least 1.0.

Based on the number of thresholds breached, firms were classified into four tiers with corresponding caps on the annual growth of interest-bearing debt: *Red* tier firms (breaching all three) faced a complete freeze on new debt; *Orange* tier firms (breaching two) were limited to 5% annual growth; *Yellow* tier firms (breaching one) to 10%; and *Green* tier firms (meeting all thresholds) to 15%. The policy was piloted with the twelve initial firms beginning in September 2020 and formally extended to the broader industry starting January 1, 2021, with a compliance deadline of end-2023.

Several features of the policy make it particularly useful for our empirical analysis. First, the announcement came with little advance warning, limiting the scope for anticipation effects. Second, the policy created sharp, well-defined variation in treatment intensity: firms' tier assignments were mechanically determined by pre-policy financial ratios, providing plausibly exogenous cross-sectional variation. Third, the policy targeted a specific sector while leaving the regulatory environment of other industries largely unchanged, allowing clean identification of spillover effects on firms outside real estate.

The consequences were severe and rapid. Many developers that entered the Red or Orange tiers—including Evergrande, Sunac, and several others—found themselves unable to roll over maturing debt or access new financing. Evergrande, which by mid-2021 carried total liabilities of approximately ¥2.4 trillion, was declared in default in December 2021. A cascade of developer defaults followed through 2022 and 2023, with widespread project suspensions, declining housing starts, and falling property prices. According to NBS data, investment in real estate development fell by roughly 10% annually in 2022–2024, and new housing starts declined by over 60% from their 2021 peak.

2.3 Financing Patterns and Trade Credit in Real Estate

Upstream transmission was possible because Chinese real estate developers financed their operations heavily through their suppliers. According to the National Bureau of Statistics, funds in place for real estate development reached RMB 19.3 trillion in 2020. Deposits and advance payments together with individual mortgage loans accounted for about half of total funds in place, self-raised funds for about one-third, and domestic loans for only 13.8%.[1] Because Chinese hous-

[1] National Bureau of Statistics, statistics on real estate development investment and sales, January–December 2020.

ing is predominantly sold before completion, that first category makes households the sector's largest single financier—and leaves them holding a claim on future delivery rather than a completed dwelling, a feature that will matter when we turn to housing demand in Section 4.2.4. These official categories, however, do not separately report trade credit—the deferred payment arrangements through which developers obtained goods and services from contractors, building-material suppliers, and other upstream firms without immediate cash payment. Yet such deferred payment was a substantial source of non-bank financing for developers.

In practice, developers routinely delayed payments to construction contractors, building-material suppliers, decoration firms, and service providers, effectively using their supply chains as a source of short-term financing. A common instrument was the commercial acceptance bill, a firm-issued promise to pay at a future date, typically with maturities of 6–12 months. Instead of receiving cash, suppliers accepted these bills from developers and could either hold them until maturity, discount them for cash, or endorse them to their own creditors, creating chains of credit that extended well beyond the initial developer-supplier relationship. Shanghai Commercial Paper Exchange data show that the year-end outstanding balance of the entire commercial bill market reached RMB 14.1 trillion in 2020, equivalent to about 14% of GDP.[2]

The reliance on supplier financing was substantial at the firm level. At year-end 2020, Evergrande reported RMB 829.2 billion in trade and other payables, representing about 43% of its total liabilities.[3] These were balance-sheet obligations, but their scale shows how heavily large developers leaned on suppliers, contractors, and other counterparties for non-bank financing. This pattern—substituting trade credit and payables for formal external financing when conventional financing becomes constrained—is precisely the transmission mechanism we formalize in our model and document empirically.

When the credit shock came in the form of the Three Red Lines policy, this pre-existing channel of supply-chain financing became visibly strained. As developer liquidity deteriorated, obligations to suppliers were no longer merely delayed payments; they became impairment risks. Listed upstream firms began reporting substantial credit losses on receivables owed by real estate clients in their 2021 disclosures, providing concrete, if anecdotal, illustrations of the transmission.

[2]Shanghai Commercial Paper Exchange, 2020 annual commercial bill market data.

[3]China Evergrande Group, Annual Report 2020. Trade and other payables: RMB 829,174 million; total liabilities: RMB 1,950,728 million.

The paint manufacturer SKSHU Paint reported RMB 814 million in credit impairment losses in 2021; at year-end 2021, its bills receivable, accounts receivable, and other receivables had a combined gross carrying amount of RMB 5.57 billion, equivalent to 45.03% of total assets.[4] The interior decoration contractor Suzhou Gold Mantis reported RMB 7.73 billion in receivables and related claims against Evergrande at year-end 2021, including RMB 2.50 billion of overdue commercial bills, RMB 1.76 billion of not-yet-due commercial bills, RMB 1.68 billion of accounts receivable, and RMB 1.80 billion of bills to be settled through asset swaps.[5] Guangtian Group provides an even starker case: it reported a 2021 net loss of RMB 5.59 billion, and its annual report explicitly attributed the decline in project activity to Evergrande's debt default.[6] At year-end 2021, Guangtian held RMB 1.83 billion in bills receivable from China Evergrande Group and its subsidiaries, against which it recorded RMB 917 million in impairment provisions, with the impairment reason stated as debt default. These disclosures, beyond their individual financial significance, illustrate that the trade-credit channel transmitted developer distress to identifiable upstream firms with real economic consequences—the dynamic our empirical analysis quantifies systematically in the following sections.

The prevalence of trade credit in real estate supply chains thus created a built-in channel for propagating financial shocks upstream. When the policy cut developers off from external financing, they leaned harder on suppliers, pushing their own liquidity stress upstream. From the suppliers' perspective, this amounted to an involuntary extension of credit to financially distressed counterparties—a dynamic that, as we show in the following sections, generated significant real effects throughout the economy. In short, a sharp and well-dated cap on the leverage of a systemically central sector landed on a supply chain that had long financed that sector through deferred payment; the remainder of the paper measures and explains what happened next.

[4]SKSHU Paint, Annual Report 2021. Credit impairment losses comprise RMB 179.0 million on bills receivable, RMB 523.3 million on accounts receivable, and RMB 111.5 million on other receivables.

[5]Suzhou Gold Mantis, 2021 performance forecast. The disclosure refers to a distressed real estate customer, widely understood to be Evergrande.

[6]Guangtian Group, Annual Report 2021.

3 Reduced-Form Evidence

This section documents how the Three Red Lines policy traveled through the production network. We describe our data, construct a firm-level measure of exposure to real estate through trade credit, and present the evidence in four steps: the economy-wide reach of the shock, the financing response of the targeted developers, the receivables and real outcomes of their suppliers, and a joint test of where in the network—and within which firms—the channel concentrates. Throughout, the design isolates variation in exposure through trade credit while holding fixed the industry-wide consequences of the housing downturn. The mechanism these facts reveal—financial distress at constrained buyers, transmitted upstream through stretched payment terms—is what the model of Section 4 then formalizes and quantifies.

3.1 Data

Our empirical analysis relies on three main sources: quarterly financial statements of listed firms, disclosed supply-chain relationships, and China's input-output table.

Firm-level financial data. We obtain quarterly data from the financial statements of Chinese listed firms in the China Stock Market & Accounting Research (CSMAR) database, covering the period from 2017Q1 to 2025Q3. Our primary variables include the components of trade credit—accounts receivable, notes receivable, advance payments, accounts payable, notes payable, and advance receipts—alongside key financial variables such as total assets, cost of goods sold, and operating revenue. We also use firms' industry classifications based on the National Industry Classification Standard (GB/T 4754). For observations with missing industry codes, we supplement using the WIND database.

Following standard practice in the literature, we exclude several categories of firms: those designated as Special Treatment (ST or ST*), B-shares (foreign-invested enterprises listed domestically), firms listed on the Beijing Stock Exchange, and financial sector firms. Our final sample consists of 5,242 listed firms from the Shanghai and Shenzhen stock exchanges, yielding 143,018 firm-quarter observations. Within this sample, 121 firms are classified in the real estate development industry, contributing 3,856 observations. One implication of the listed-firm sample is worth stating at the outset: listing requires scale and financial strength, so the suppliers we observe are

among the firms best placed to absorb payment delays. To the extent that smaller, unlisted suppliers bear more of the channel, our firm-level estimates describe its incidence among the most resilient firms.

Supply-chain data. To identify direct customer–supplier relationships, we use CSMAR's supply-chain database, which reports each listed firm's top five customers and top five suppliers as disclosed in annual reports. While this dataset has limitations—it covers only disclosed relationships and is available for a subset of firms—it represents the only source of micro-level supply-chain linkages in China. We use these disclosed relationships to construct treatment indicators for firms directly connected to the real estate sector.

Input-output table. We obtain the 2020 input-output table from China's National Bureau of Statistics. This table provides industry-level data on flows of intermediate goods across 153 sectors, from which we compute the Leontief inverse matrix that captures both direct and indirect production linkages. We use this matrix to construct measures of treatment intensity based on each industry's exposure to real estate through the production network.

Table 1 reports summary statistics. Together, the three sources let us observe both sides of the trade-credit relationship—payables at developers, receivables at their suppliers—and each firm's position in the production network.

3.2 Treatment Intensity and Specification

Studying how a financial shock to real estate propagates through the economy requires a measure of each firm's exposure to that shock through trade credit. In the absence of bilateral data on trade credit, we construct a measure of treatment intensity that combines firm-level information on reliance on trade credit with industry-level linkages in the production network.

Specifically, for each firm i in industry k, we define:

$$\text{Intensity}_i = \underbrace{\overline{\left(\frac{\text{AR}_{it}}{\text{Assets}_{it}}\right)}_{t \le 2020\text{Q2}}}_{\text{Firm-level trade credit reliance}} \times \underbrace{\Psi_{kh}}_{\text{Industry exposure to RE}} \tag{1}$$

Table 1: Summary statistics

	N	Mean	SD	Min	Max
Panel A. Treatment variables (firm-level, $t \leq$ 2020Q3)					
AR / Assets	5,182	0.1343	0.0984	0.0000	0.7291
Leontief exposure $\Psi_{i,h}$	5,182	0.0663	0.1571	0.0029	1.0627
Intensity (pp)	5,182	0.6869	1.3064	0.0000	35.5747
Panel B. Outcome variables, full sample (firm-quarter, 100M RMB)					
Cash	142,182	26.28	117.61	0.00	4,243.96
Sales (quarterly)	139,603	31.34	204.33	0.00	8,762.59
Total assets	142,182	197.75	997.62	0.49	34,125.21
Interest-bearing debt	142,182	34.73	179.90	0.00	6,998.60
Accounts payable	142,182	23.05	177.05	0.00	9,482.32
Accounts receivable	142,182	16.25	76.15	0.00	4,132.08
Panel C. Real estate firms only (K70)					
Accounts payable	3,856	88.82	349.83	0.00	4,101.24
AP / Sales	3,832	2.555	8.124	0.000	419.853
Cash	3,856	93.24	212.74	0.04	1,965.99
Short-term loans	3,856	19.54	45.71	0.00	400.52

Notes. Panel A reports the cross-sectional distribution of treatment variables used in the firm-level intensity construction. AR/Assets is the firm-level pre-policy average ($t \leq$ 2020Q3) of accounts receivable scaled by total assets; $\Psi_{i,h}$ is the Leontief inverse element from industry i to the real estate sector h in the 2020 China I-O table; Intensity is the firm-level interaction scaled by 100. Panel B reports the firm-quarter distribution of outcome variables across the full sample (Shanghai-Shenzhen listed firms, 2017Q1–2025Q3, financial firms and ST firms excluded). Panel C restricts to firms in the K70 (real estate development) industry; AP/Sales uses quarterly sales as the denominator. Monetary values are in 100M RMB unless otherwise noted.

The first component is the firm's average ratio of accounts receivable to total assets during the pre-policy period, capturing how much trade credit the firm extends to its customers, and hence how vulnerable it is when those customers delay payment. Importantly, this component is pre-determined: it is calculated entirely from data before the policy announcement and is therefore not affected by the policy itself. We verify that this ratio exhibits no pre-policy trend ($p = 0.47$ for a time trend in a firm fixed effects regression restricted to the pre-policy sample), confirming that our measure captures a stable firm characteristic rather than a pre-existing trajectory.

The second component is the element of the Leontief inverse matrix $\Psi = (\mathbb{I} - \Omega)^{-1}$ corresponding to industry k's total (direct and indirect) input requirements from the real estate sector h. This captures the extent to which an industry relies on real estate as a supplier through the production network, accounting for higher-order linkages.

The interaction of these two components creates firm-level variation in exposure *within* industries: even among firms in the same industry (and thus facing the same IO linkage to real estate),

those with higher pre-policy AR ratios are more exposed to disruptions in trade credit. This is the identifying variation that our empirical strategy exploits. We scale the intensity measure by 100 so that coefficients can be interpreted as the effect of a one-percentage-point increase in exposure.

Specification. Our main specification is:

$$Y_{it} = \beta \cdot (\text{Intensity}_i \times \text{Post}_t) + \theta_i + \gamma_{kt} + \varepsilon_{it} \tag{2}$$

where Y_{it} is an outcome variable for firm i in quarter t, Post_t equals one for $t \geq$ 2020Q3 and zero otherwise, θ_i denotes firm fixed effects, and γ_{kt} denotes industry-by-time fixed effects. The inclusion of industry-by-time fixed effects is important: it absorbs any time-varying shocks common to firms within the same industry, ensuring that our estimates are identified purely from within-industry, across-firm variation in exposure through trade credit. Standard errors are clustered at the industry level.

Because the industry component of the exposure measure is absorbed by the industry-by-time fixed effects, identification rests on the within-industry AR/Assets gradient. We therefore report, alongside the baseline, a specification in which pre-policy AR/Assets and pre-policy size each carry their own post-policy trajectory, so that the network interaction is not identified off either characteristic on its own (Panel C of Table 2).

3.3 Economy-Wide Spillover Effects

We begin with our central reduced-form finding: the Three Red Lines policy generated spillovers that reached well beyond the targeted real estate sector, in proportion to firms' exposure through trade credit. Table 2 reports estimates of equation (2) for four outcomes—cash holdings, quarterly sales, interest-bearing debt, and equity—each measured in 100M RMB. We deliberately exclude total assets from the outcome list: total assets is the denominator of the AR/Assets component of the exposure measure, so an estimated effect on assets would be partly mechanical.

Panel A is our baseline. The estimates indicate broad-based economic deterioration among firms more exposed to real estate through trade credit. A one-percentage-point increase in exposure is associated with a decline of 236 million RMB in cash holdings (column 1), 204 million RMB

Table 2: Economy-Wide Spillover Effects

Panel A: baseline	Dependent Variable (100M RMB)			
	(1) Cash	(2) Sales	(3) Debt	(4) Equity
Intensity(%) × Post	-2.36***	-2.04**	-2.25***	-3.62***
	(0.54)	(0.80)	(0.69)	(0.84)
N	142024	139429	142024	142024
Firm FE	Yes	Yes	Yes	Yes
Industry × Time FE	Yes	Yes	Yes	Yes
Pre-policy size × Post	No	No	No	No
Pre-policy AR × Post	No	No	No	No
Adj. R^2	0.899	0.951	0.897	0.976
Panel B: with pre-policy size control				
Intensity(%) × Post	-3.53***	-3.42***	-3.80***	-5.78***
	(0.61)	(1.00)	(1.07)	(0.89)
Pre-policy emp × Post	14.33***	19.64***	18.99**	29.79***
	(3.09)	(5.99)	(9.07)	(6.11)
N	114771	113783	114771	114771
Firm FE	Yes	Yes	Yes	Yes
Industry × Time FE	Yes	Yes	Yes	Yes
Pre-policy size × Post	Yes	Yes	Yes	Yes
Pre-policy AR × Post	No	No	No	No
Adj. R^2	0.899	0.953	0.898	0.973
Panel C: both predetermined characteristics × Post				
Intensity(%) × Post	-2.38***	-2.55***	-2.08*	-3.54**
	(0.61)	(0.67)	(1.09)	(1.42)
Pre-policy AR/Assets(%) × Post	-0.31**	-0.23	-0.47*	-0.61***
	(0.14)	(0.15)	(0.25)	(0.21)
Pre-policy emp × Post	14.22***	19.56***	18.83**	29.58***
	(3.05)	(5.96)	(8.98)	(6.04)
N	114771	113783	114771	114771
Firm FE	Yes	Yes	Yes	Yes
Industry × Time FE	Yes	Yes	Yes	Yes
Pre-policy size × Post	Yes	Yes	Yes	Yes
Pre-policy AR × Post	Yes	Yes	Yes	Yes
Adj. R^2	0.899	0.953	0.898	0.973

Notes. Each column reports a separate estimate of equation (2) on the sample of Shanghai–Shenzhen listed firms. Outcomes are in 100M RMB. Intensity is the firm's pre-policy mean AR/Assets multiplied by its industry's Leontief exposure to real estate, scaled by 100. Panel B adds the interaction of pre-policy log employment with Post and accordingly restricts the sample to firms with non-missing pre-policy employment. Panel C adds, alongside it, the interaction of pre-policy AR/Assets with Post, so that both predetermined firm characteristics are allowed their own post-policy trajectory and the network interaction is identified off neither one alone. Standard errors, clustered at the industry level, in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

in quarterly sales (column 2), 225 million RMB in interest-bearing debt (column 3), and 362 million RMB in equity (column 4). All four coefficients are statistically significant at the 5% level or better.

The pattern across outcomes tells a coherent story of financial contagion. The cash decline indicates a direct liquidity squeeze: firms exposed to real estate experienced a drain on their liquid resources, consistent with delayed receipt of payments from customers tied to real estate. The sales decline suggests that the financial shock translated into a contraction of real activity—exposed firms either lost demand from the contracting real estate sector or scaled back their own operations as liquidity tightened. The simultaneous decline in both debt and equity implies a deterioration in firms' overall financial position, not merely a temporary disruption of cash flow.

Panel B asks whether the baseline merely reflects firm size: if high-exposure firms are systematically larger or smaller, differential size trends after 2020 could masquerade as a trade-credit effect. The control must be predetermined: contemporaneous employment is itself an outcome of the shock—conditioning on it would be a classic bad control—and it is collinear with the size-like outcomes we study, whereas the interaction of pre-policy log employment with Post_t is subject to neither objection. Adding this predetermined size control strengthens rather than weakens the estimates: the implied declines per percentage point of exposure rise to 353 million RMB for cash, 342 million for sales, 380 million for debt, and 578 million for equity, all significant at the 1% level. We treat Panel A as the baseline specification—it is also the conservative input for the comparison with the model in Section 5—and read Panel B as evidence that differential size trends do not explain the result.

Panel C confronts the sharper version of the same concern. Because Ψ_{kh} is constant within an industry and industry-by-time fixed effects absorb every industry-level movement, the surviving variation in the interaction is the within-industry gradient in pre-policy AR/Assets, scaled by the industry's exposure. A reader may therefore ask whether the estimate reflects nothing more than the fragility of trade-credit-reliant firms in general: firms carrying large receivables in 2019 might have fared worse after 2020 for reasons—the pandemic, the working-capital tightening of 2021–22—that require no role for the real estate network. Panel C answers it directly by adding the interaction of pre-policy AR/Assets with Post_t alongside the size control, so that both predetermined characteristics carry their own post-policy trajectory and the network interaction must earn its coefficient from the residual variation. It does: the estimates are -2.38 for cash, -2.55 for

sales, -2.08 for debt, and -3.54 for equity, statistically significant and close to the Panel A baseline. The general AR effect is real—its own coefficient is negative and significant in three of four columns—but it does not account for the network gradient.

The two controls must enter together, and the reason is instructive. Pre-policy AR/Assets and firm size are entangled within industries, and they bias the interaction in opposite directions: adding the AR interaction alone attenuates it (to -1.02 for cash, statistically indistinguishable from zero) because AR/Assets partly proxies for size, while adding the size interaction alone strengthens it (to -3.53). Only with both in the specification is the network interaction identified off neither one. Allowing each characteristic an unrestricted quarterly path rather than a single post-policy shift—replacing the two interactions with a full set of characteristic-by-quarter terms—leaves the estimates essentially unchanged (-2.36 for cash, -2.53 for sales), so the linear specification in Panel C is not doing the work.

Dynamic effects. Figure 1 presents event study estimates that trace out the time path of the cash effect. Two features stand out. First, coefficients in the pre-policy quarters are close to zero and statistically insignificant, supporting the identifying assumption that more-exposed and less-exposed firms within the same industry were on parallel trajectories before the policy. Second, the decline in cash holdings emerges immediately after the policy and deepens over time, with no sign of recovery through the end of our sample. The persistence of this effect—lasting over three years—suggests that the trade-credit shock set off a sustained process of financial deterioration rather than a transitory adjustment.

Interpretation of the continuous slope. Callaway et al. (2024) note that the continuous-treatment difference-in-differences slope identifies a weighted average of dose-specific level effects under standard parallel trends, but recovering a marginal causal response to dose requires the stronger assumption that high-exposure firms would have responded identically to lower exposure. We therefore read β in the main specification as a parsimonious summary of how post-policy outcomes vary with pre-determined exposure rather than as a structural marginal effect. Appendix A.3 reports a quintile dose-response and binary high-versus-low splits at the quartile, tercile, and median cutoffs that together confirm the underlying fact: firms more exposed ex ante to real estate

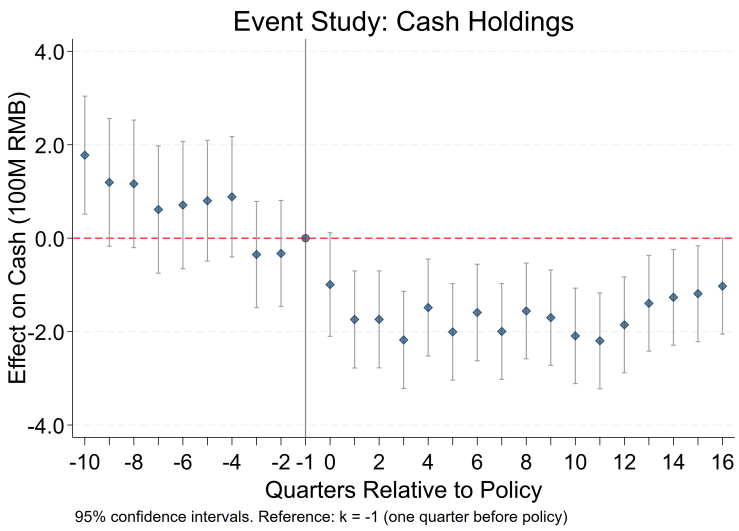


Figure 1: Event Study: Cash Holdings

Notes: This figure plots the coefficients β_k from the event study regression of cash holdings (100M RMB) on Intensity (%) interacted with period dummies. The reference period is one quarter before the policy (2020Q2). The vertical line marks the policy quarter (2020Q3). Bars represent 95% confidence intervals.

through trade-credit channels deteriorate more after the policy, with the lowest-exposure quintile statistically indistinguishable from the reference and a monotone gradient through the upper four.

These results establish that the Three Red Lines policy had consequences far beyond the real estate sector itself: firms throughout the economy lost liquidity, revenue, and balance-sheet strength in proportion to their exposure to real estate through trade credit. A natural question follows: through what mechanism did this transmission occur? We answer it from both ends of the credit relationship, examining first the shocked sector itself and then its upstream suppliers.

3.4 The Shocked Sector: Real Estate Developers

The trade-credit channel can operate only if the policy in fact pushed real estate firms to delay payments to their suppliers. This subsection provides direct evidence on the "sending end" of the transmission: whether, as the policy withdrew bank credit, developers substituted away from prepayment—so that accounts payable rose while short-term bank loans and cash holdings fell.

3.4.1 Cross-Industry Evidence

We first compare real estate firms to all other listed companies using a standard difference-in-differences design:

$$\begin{aligned} Y_{it} = \beta \cdot (\mathrm{RE}_i \times \mathrm{Post}_t) + \delta \cdot (\mathrm{Emp}_i^{\mathrm{pre}} \times \mathrm{Post}_t) \\ + \theta_i + \lambda_t + \varepsilon_{it} \end{aligned} \tag{3}$$

where RE_i indicates real estate development firms, Post_t equals one from 2020Q3 onward, and $\mathrm{Emp}_i^{\mathrm{pre}}$ is the firm's pre-policy log employment, the predetermined size control introduced above. The specification includes firm and quarter fixed effects, with standard errors clustered at the industry level. The debt outcome here is short-term bank loans—precisely the credit margin the policy cut.

Table 3 reports the results. Short-term bank loans contracted by 678 million RMB per developer (column 1), and cash holdings fell by 500 million RMB (column 2). Accounts payable, in contrast, surged by 5.58 billion RMB (column 3). The picture is one of substitution: as bank credit dried up and internal liquidity drained, real estate firms shifted the financing burden onto their supply-chain partners. From the perspective of suppliers, this amounted to an involuntary extension of credit to financially distressed counterparties. The dynamics behind these averages are reported in Appendix F: Figure 11 shows that the rise in accounts payable emerges sharply after 2020Q3 and accumulates over the following quarters, while Figure 12 shows a contemporaneous and persistent decline in short-term bank loans—consistent with the substitution from bank credit to supplier financing that the static estimates imply.

3.4.2 Within-Sector Variation

The cross-industry design, while informative, sets real estate firms against a heterogeneous set of other industries. To sharpen identification, we exploit the policy's tiered structure and compare firms within the real estate sector according to whether they breached at least one of the three regulatory thresholds (the Red, Orange, or Yellow tiers) as of 2020Q2, the last pre-policy quarter,

Table 3: The Shocked Sector: Real Estate Developers versus Other Listed Firms

	Dependent Variable (100M RMB)		
	(1) ST Loan	(2) Cash	(3) AP
RE × Post	-6.78***	-5.00***	55.82***
	(0.43)	(0.80)	(1.69)
Pre-policy emp × Post	1.96	14.91***	27.33**
	(1.89)	(3.13)	(11.59)
N	114966	114966	114966
Firm FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
Pre-emp control	Yes	Yes	Yes
Adj. R^2	0.819	0.895	0.847

Notes. RE indicates real estate development (K70) firms; Post equals one from 2020Q3. ST Loan is short-term bank loans; AP is accounts payable; outcomes are in 100M RMB. All columns include firm and quarter fixed effects and the interaction of pre-policy log employment with Post. Standard errors, clustered at the industry level, in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

versus meeting all three (the Green tier):

$$Y_{it} = \beta \cdot (\text{Breached}_i \times \text{Post}_t) + \theta_i + \lambda_t + \varepsilon_{it} \tag{4}$$

This specification uses only real estate firms and clusters standard errors at the firm level. It deliberately includes no size control: breaching is itself a size and leverage characteristic of the developer, so interacting size with Post would absorb the very treatment we want to measure.

Table 4 shows that the response was concentrated among the firms most directly bound by the policy. Breached developers increased accounts payable by 8.71 billion RMB relative to green-tier developers (column 3) and cut short-term borrowing by 1.01 billion RMB (column 1); the difference in cash holdings is negative but imprecisely estimated (column 2). Figure 13 in Appendix F traces the dynamic response of accounts payable for breached firms relative to green-tier developers, confirming that the divergence emerges only after the policy and widens steadily through the post-period.

Together, the evidence from the shocked sector establishes the first link in the transmission chain: the Three Red Lines policy cut developers' access to short-term bank credit, and developers leaned on supplier financing in response. The accounts payable of real estate firms—which are the accounts receivable of their suppliers—surged, transmitting the liquidity shock upstream through

Table 4: Within Real Estate: Breached versus Green-Tier Developers

	Dependent Variable (100M RMB)		
	(1) ST Loan	(2) Cash	(3) AP
Breached × Post	-10.05**	-3.34	87.12**
	(4.47)	(12.25)	(34.07)
N	3856	3856	3856
Firm FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
Pre-emp control	No	No	No
Adj. R^2	0.740	0.893	0.825

Notes. Sample: real estate development (K70) firms only. Breached equals one for developers in violation of at least one of the three regulatory thresholds as of 2020Q2. ST Loan is short-term bank loans; AP is accounts payable; outcomes are in 100M RMB. All columns include firm and quarter fixed effects; no size control is included (see text). Standard errors, clustered at the firm level, in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

the production network.

3.5 The Receiving End: Upstream Suppliers

That developers delayed payments does not by itself establish that suppliers were harmed—suppliers might have held buffers thick enough to absorb the delayed receipts. We now examine the receiving end of the trade-credit relationship, asking whether upstream suppliers experienced rising receivables and real deterioration. We use two complementary designs, and the trade-credit mechanism tells us in advance where each signature should appear. The first design observes the seller–buyer tie directly, so the receivables response should be visible there. In the second design, exposure varies only at the industry level; it covers the whole economy but dilutes any firm-specific receivables signal, so it should instead pick up the broader contraction of liquidity and credit.

3.5.1 Evidence from Disclosed Customer Relationships

We exploit micro-level supply-chain data from CSMAR, which reports each listed firm's top five customers as disclosed in annual reports. A firm is defined as treated if it disclosed at least one real estate customer during 2018–2020, before the policy. We compare these firms to those without disclosed real estate customers in a difference-in-differences framework using annual data, with

Post_t equal to one for $t \geq 2021$ and zero otherwise:

$$\begin{aligned} Y_{it} = \beta \cdot (\text{Treat}_i \times \text{Post}_t) + \delta \cdot (\text{Emp}_i^{\text{pre}} \times \text{Post}_t) \\ + \theta_i + \gamma_t + \varepsilon_{it} \end{aligned} \tag{5}$$

To ensure comparability, we restrict the sample to a balanced panel of firms observed both before and after the policy. The specification includes firm and year fixed effects together with the predetermined size control, and clusters standard errors at the industry level.

Table 5 presents the results, and the receivables prediction is borne out where the customer tie is observed. The ratio of accounts receivable to sales rises by 0.115 among treated suppliers after the policy (column 1)—receivables climb by 11.5 percentage points of sales—as developers' payment delays land on their suppliers' balance sheets. The shock does not stop there: log sales fall by 0.229 (column 4), so treated suppliers also contract in real terms. The cash-to-assets and debt-to-assets ratios (columns 2 and 3) carry negative point estimates but are statistically insignificant in this small annual panel. Figure 14 in Appendix F reports the corresponding event study for the AR/Sales ratio, showing flat pre-trends followed by a sustained rise in receivables relative to sales after 2020.

One margin of adjustment deserves note. A supplier that saw developer distress coming could sever the relationship—and to the extent any did, exit works against our finding: a supplier that stops selling to developers accumulates no new receivables, so anticipatory severance pushes the estimated receivables response toward zero. The 11.5-percentage-point rise is net of any such exit.

3.5.2 Evidence from Input-Output Linkages

The disclosed-customer analysis is sharp but narrow: it covers only the subset of firms that voluntarily report customer identities. To assess whether the transmission operated at scale, we exploit industry-level input-output linkages across the full sample of non-real-estate listed firms:

$$\begin{aligned} Y_{it} = \beta \cdot (\Psi_{kh} \times \text{Post}_t) + \delta \cdot (\text{Emp}_i^{\text{pre}} \times \text{Post}_t) \\ + \theta_i + \lambda_t + \varepsilon_{it} \end{aligned} \tag{6}$$

Table 5: Suppliers with Disclosed Real Estate Customers

	(1) AR/Sales	(2) Cash/Assets	(3) Debt/Assets	(4) ln(Sales)
Treat × Post	0.115**	-0.026	-0.006	-0.229**
	(0.054)	(0.020)	(0.028)	(0.108)
Pre-policy emp × Post	0.004	0.008**	-0.008	-0.024
	(0.009)	(0.004)	(0.005)	(0.029)
N	2301	2276	1851	2320
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Pre-emp control	Yes	Yes	Yes	Yes
Adj. R^2	0.486	0.690	0.751	0.938

Notes. Annual balanced panel of listed firms with disclosed top-five customers. Treat equals one for firms disclosing at least one real estate customer during 2018–2020; Post equals one from 2021. AR/Sales, Cash/Assets, and Debt/Assets are ratios; ln(Sales) is log sales. All columns include firm and year fixed effects and the interaction of pre-policy log employment with Post. Standard errors, clustered at the industry level, in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

where Ψ_{kh} is the Leontief inverse element capturing industry k's total production linkage to the real estate sector, scaled by 100 so that coefficients are interpreted per percentage point of exposure.

Two features of this design are deliberate. First, the treatment is the Leontief exposure alone—not the interaction with pre-policy AR/Assets used in equation (2)—because accounts receivable is now the outcome: weighting the treatment by a firm's own receivables would build the dependent variable into the regressor and render the receivables response mechanically circular. Second, because the treatment varies only at the industry level, industry-by-time fixed effects would absorb it; identification therefore relies on firm and quarter fixed effects, with standard errors clustered at the industry level and the predetermined size control included.

Table 6 reports the results. The broad design picks up the financial contraction: per percentage point of exposure, cash falls by 70 million RMB (column 2) and debt by 123 million RMB (column 3). The receivables coefficient (column 1) is small and statistically indistinguishable from zero, and the sales coefficient (column 4), though negative, is also imprecise—exactly what one expects when an industry-level exposure measure is asked to detect a response that varies firm by firm with actual customer ties.

The joint pattern also speaks to an alternative reading in which banks, observing receivables pile up on suppliers' balance sheets, simply redirect credit toward better uses—a reallocation that

would improve the allocation of financial resources rather than transmit distress. Under that reading, displaced firms should lose bank credit while real activity holds. Instead, exposed firms cut capital expenditure, run down inventories, and earn lower gross profit alongside the joint decline in cash and debt (Appendix A.5). Financing and real activity contract together—the signature of a tightening, not of an efficient redirection of credit.

Table 6: Suppliers: Exposure through Input-Output Linkages

	Dependent Variable (100M RMB)			
	(1) AR	(2) Cash	(3) Debt	(4) Sales
RE exposure × Post	-0.22	-0.70*	-1.23**	-0.65
	(0.25)	(0.40)	(0.59)	(0.66)
Pre-policy emp × Post	9.31**	15.84***	18.52**	20.05***
	(4.12)	(3.12)	(9.29)	(6.22)
N	111510	111510	111510	110535
Firm FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Pre-emp control	Yes	Yes	Yes	Yes
Adj. R^2	0.862	0.895	0.883	0.949

Notes. Sample: listed firms outside real estate. RE exposure is the industry's Leontief linkage to the real estate sector, scaled by 100. AR is accounts receivable; outcomes are in 100M RMB. All columns include firm and quarter fixed effects and the interaction of pre-policy log employment with Post. Standard errors, clustered at the industry level, in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Taken together, the two designs are complementary in precisely the way a trade-credit channel predicts. The sharp, firm-level identification—where we observe who actually sells to developers—carries the receivables and sales response; the broad, industry-level exposure—which covers the whole economy but cannot see individual customer ties—carries the contraction of liquidity and debt. The receivables signature appears only where the seller–buyer tie is observed directly. A generic demand shock to industries linked to real estate would give no particular reason for the receivables response to concentrate in the disclosed sample; a channel operating through payment terms on actual supplier relationships predicts exactly that concentration. The supply-side evidence thus validates the mechanism from both ends of the credit relationship: the payables that surged at developers (Section 3.4) reappear as receivables, lost liquidity, and lost sales at their suppliers.

Mechanism dynamics, heterogeneity, and real-economy outcomes (Appendix A). The supply-side findings extend in several directions that we report in detail in Appendix A. First, the channel shows up directly in trade-credit metrics: more exposed firms experience significantly longer days-sales-outstanding and higher accounts-receivable-to-sales ratios after the policy, consistent with the payment-delay accumulation that drives the channel (Appendix A.4). Second, the financial shock propagates into real production decisions: high-intensity firms cut capital expenditure, run down inventory, and earn lower gross profit, evidence that the channel does not stop at firm balance sheets but constrains real activity (Appendix A.5). Third, the absorption is highly uneven across firms: private and smaller firms—those with limited access to substitute financing—bear nearly all of the channel's bite, while state-owned and larger firms either remain largely insulated or substitute toward bank borrowing to bridge the shock (Appendix A.6). Finally, a continuous dose-response specification confirms a monotone gradient across the intensity distribution (Appendix A.3). Together these auxiliary results already suggest that the channel is fundamentally financial—operating through firms' working-capital positions, and concentrated where alternative financing margins are unavailable. The next subsection develops this interpretation by asking, more sharply, *where* in the production network and *which* firms absorb the shock, and by linking the empirical signature directly to the model's microstructural assumption.

3.6 Network Position and Financial Constraint

Sections 3.4–3.5 establish the existence of the trade-credit channel and verify it from both ends of the credit relationship. We now sharpen the analysis with two structural questions: *where* in the production network does the channel propagate, and *within whom* does it concentrate.

We decompose each industry's Leontief exposure to real estate into a direct component A_{ih} (the direct input share to real estate) and a higher-order residual $(A^2 + A^3 + \cdots)_{ih}$ (transmission through multi-step paths $i \to k \to h$, $i \to k \to l \to h$, and so on). On the constraint dimension, we partition firms by ownership as of January 1, 2020, treating private firms as more constrained than state-owned enterprises (SOEs)—a distinction that yields consistent results when we instead proxy constraint by firm size (Appendix A.6). Table 7 estimates the channel by network order, by ownership, and by their interaction on the standard cash outcome.

Table 7: Network Position and Financial Constraint — Joint Test (Cash, 100M RMB)

	Cash (100M RMB)				
	(1) Higher	(2) Higher×	(3) Direct	(4) Direct×	(5) Both
Higher × Post [SOE baseline]	-2.948***	-2.498***			-2.402***
	(0.294)	(0.522)			(0.512)
Higher × Post × Private		-1.143**			-1.372***
		(0.558)			(0.446)
Direct × Post [SOE baseline]			-1.785	-2.724	-0.443
			(1.693)	(2.533)	(0.574)
Direct × Post × Private				1.414	0.988
				(1.146)	(0.788)
N	142024	103965	142024	103965	103965
Adj. R^2	0.899	0.895	0.899	0.895	0.895

Private firms = (Wind 2020-01-01 snapshot). Baseline = SOE.
Cols. 2,4,5 restrict to SOE+Private (other categories dropped).
Firm + Time + Industry-by-Time FE. Cluster SE at industry.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

The pattern in Table 7 is sharply structured. The channel runs almost entirely through higher-order network paths: direct exposure to real estate produces no detectable effect, whether in isolation, jointly with higher-order terms, or within either ownership group, while higher-order exposure carries a substantial and robust effect throughout. Direct sellers to real estate—dominated in China by finance, business services, and utilities—do not transmit the shock upstream, plausibly because their alternative liquidity margins absorb the demand contraction without forcing them to squeeze their own suppliers. Transmission through trade credit is, at bottom, a relay: it operates where each intermediate firm is itself squeezed hard enough to stretch its own payables, and the multi-step manufacturing chains that the higher-order component captures are where that relay runs. Nor is the comparison between the two components mechanical: both are measured in percentage points of Leontief-weighted exposure, so the coefficients compare effects *per unit* of exposure—that the higher-order component aggregates many paths affects its level, not its per-unit coefficient. Within the chain of higher-order suppliers that does carry the channel, financial constraint deepens the hit a firm takes: private firms suffer roughly half again the cash decline of SOEs at the same exposure, consistent with their lack of substitute financing. Neither network

position nor constraint is sufficient on its own—the effect emerges multiplicatively, only where the chain reaches *and* the receiving firm cannot substitute around the involuntary lending.

This signature—indirect chain propagation interacting with firm-level constraint—adds structural detail to the findings of Sections 3.4–3.5. The trade-credit channel does not propagate uniformly through every direct buyer–supplier tie; it operates along the indirect chains that reach financially exposed firms. We return to the theoretical implications of this joint signature—and to the contrast with one-dimensional alternatives—when we build the multi-sector model in Section 4.

3.7 Robustness

We close the empirical analysis with a battery of checks on the economy-wide estimates. All supporting tables and figures are in Appendix A; the text here reports only what we do and what we find.

First, size. The most natural confound is that exposure proxies for firm size, with large and small firms simply trending differently after 2020. Panel B of Table 2 addresses this with the predetermined size control—the interaction of pre-policy log employment with Post—and the four baseline effects survive and in fact strengthen: the cash effect moves from 236 to 353 million RMB per percentage point of exposure, and the sales effect from 204 to 342 million. The designs of Sections 3.4 and 3.5 include the same control throughout, with one deliberate exception: the within-developer comparison, where breaching is itself a size and leverage characteristic and a size-by-Post control would absorb the treatment. The appendix results retain the parsimonious Panel A specification, which is also the conservative choice for the comparison with the model in Section 5, since the size-controlled coefficients are uniformly larger in magnitude. Appendix A collects the supporting results on trade-credit dynamics, real outcomes, and heterogeneity.

Second, functional form. The continuous-treatment slope in equation (2) is not an artifact of linearity. Appendix A.3 reports a quintile dose-response together with binary high-versus-low splits at the quartile, tercile, and median cutoffs: the lowest-exposure quintile is statistically indistinguishable from the reference group, the gradient is monotone through the upper four quintiles, and the binary splits deliver the same qualitative answer at every cutoff. The post-policy deterio-

ration rises with ex-ante exposure however the exposure distribution is cut.

Third, differential sensitivity to aggregate financial conditions. The industry-by-time fixed effects absorb the common component of every aggregate shock, but not its differential incidence (Donaldson, 2026): if higher-exposure firms are systematically more sensitive to the monetary and credit conditions that swung sharply over 2020–2022—the interbank rate fell by roughly a percentage point, credit growth accelerated, and the reserve requirement was cut repeatedly—then the estimates could reflect that differential sensitivity rather than the trade-credit channel. We stress-test the estimate against that loading directly, rather than by adding an ever-growing list of controls. The sharpest test uses a high-frequency, exogenous monetary policy shock (He et al., 2026); we also re-estimate equation (2) adding, one at a time, the interaction of Intensity with each of eight measures of aggregate conditions spanning prices (money-market and loan rates) and quantities (credit growth and the reserve requirement). The coefficient of interest stays negative and significant throughout and for both outcomes: the sales effect is essentially unchanged, and the cash effect, though it attenuates modestly under the rate and reserve measures, never falls below roughly 160 million RMB per percentage point of exposure. That residual attenuation is itself partly a size effect: adding the predetermined employment control shrinks it. Appendix A.1 collects the full set of results. The spillover is not an artifact of the monetary easing that accompanied the deleveraging.

Fourth, influence and inference. No single sector drives the economy-wide estimate: re-estimating with one industry dropped at a time leaves the cash and sales coefficients tightly clustered around their full-sample values and statistically significant in all but at most one of the seventy-six leave-one-out fits. And because a common aggregate shock can correlate outcomes across firms within a quarter, we recompute the standard errors clustering simultaneously by industry and by quarter; they are almost identical to the industry-clustered baseline and leave every economy-wide coefficient significant. Appendix A.2 reports both checks.

4 A Multi-Sector Model with Endogenous Trade Credit

The empirical analysis of Section 3 establishes three facts about how the Three Red Lines policy traveled through the production network. First, listed firms more exposed to real estate through

trade credit deteriorated broadly after the policy: cash, sales, debt, and equity all declined, and the declines persist for more than three years. Second, the constrained developers themselves substituted supplier financing for the bank credit the policy withdrew: accounts payable rose by 5.58 billion RMB at the average developer—and by 8.71 billion RMB among developers that breached the policy's thresholds—while their short-term bank loans and cash holdings fell. Third, the suppliers on the other end of those invoices absorbed the pressure: at firms with disclosed real-estate customers, receivables climbed by 11.5 percentage points of sales and sales contracted.

These facts identify the channel; they do not measure its aggregate cost. In the data we observe rising payables at developers and rising receivables at their suppliers, and the designs of Section 3 locate the channel's incidence—which firms absorb the pressure, and along which network paths it travels. Two objects, however, lie beyond the reach of the reduced form. The first is the counterfactual aggregate loss: the industry-by-time fixed effects that secure identification also absorb the economy-wide adjustment, so the estimates compare more-exposed with less-exposed firms and are silent on how much output the economy lost in total. The second is the decomposition: any contraction of real-estate spending propagates upstream mechanically through input-output linkages, and no regression can separate this Leontief arithmetic from the additional contraction that arises because payment terms themselves tighten. Recovering both objects is the task of the model we develop in this section.

We proceed in two steps. We first present a deliberately simple two-sector model that isolates the transmission mechanism and states it as a proposition (Section 4.1). We then embed the same prepayment choice in a seven-sector economy calibrated to China's 2020 input-output table (Section 4.2), where the prepayment condition reappears, sector by sector, as each firm's first-order condition inside the production network. Real estate enters with a structure dictated by its product: new housing is a durable asset that households accumulate, not a flow good that other sectors consume, and household demand for housing varies endogenously with credit conditions through a time-varying preference wedge. Neither feature alters the trade-credit mechanism, which operates exactly as in the two-sector model. Section 4.3 describes the calibration and Section 4.4 reports the baseline impulse response; Section 5 then puts the model to work—decomposing the trade-credit channel, confronting the model with the reduced-form estimates, locating the loss across sectors, and evaluating policy counterfactuals.

4.1 A Simple Framework

Having documented the transmission empirically, we now formalize it. We begin with a minimal model of the relationship at the heart of the paper: a borrowing-constrained downstream firm, interpretable as a real estate developer, and the supplier that finances part of its production. The model is static and partial-equilibrium—we take factor prices, the intermediate-good price, the downstream output price, and interest rates as given—because its sole purpose is to show, analytically, how a shock to one sector's borrowing limit reaches that sector's supplier through endogenous payment terms. It delivers the object our regressions estimate: the response of prepayment to borrowing tightness, $d\phi/d\theta_h$, together with the signs this response imposes on payables, bank credit, and cash at the constrained buyer and on receivables and output at its supplier. The general-equilibrium feedbacks that the framework sets aside—price adjustment and demand contraction across many interconnected sectors—are the task of the multi-sector model in Section 4.2 below.

4.1.1 Setup

There are two competitive sectors. A *downstream* sector (denoted h, interpretable as real estate) buys an intermediate good and faces a borrowing constraint; an *upstream* sector (denoted n) produces the intermediate good using labor. A representative household supplies labor and consumes; because all prices are taken as given, we do not model it explicitly.

Downstream sector. The downstream firm produces output $IH = A_h X^{\alpha}$ $(0 < \alpha < 1)$ from the intermediate input X, and sells at the price P_h. It purchases X at unit price P, prepaying a fraction $\phi \in [0,1]$ of the bill in cash and deferring the remainder $(1-\phi)$ as trade credit. Prepayment is financed by bank borrowing at gross cost R_h; deferral carries a convex trade-credit cost $G(\phi)\,PX$ with $G(\phi) = \frac{1}{2}(1-\phi)^2$, so that relying more heavily on trade credit (lower ϕ) is more costly at the margin. The firm's bank-financed expenditure is limited by a borrowing constraint tied to the value of its sales:

$$\phi\, PX \;\leq\; \theta_h\, P_h\, IH, \tag{7}$$

where θ_h is the (shock-able) borrowing tightness. The firm chooses X and ϕ to maximize profit

$$\Pi_h \;=\; P_h A_h X^{\alpha} \;-\; PX \;-\; G(\phi)\, PX \;-\; R_h\, \phi\, PX$$

subject to (7). Letting $\lambda_h \geq 0$ denote the multiplier on (7), the first-order condition for ϕ is $-G'(\phi) = R_h + \lambda_h$, i.e.

$$\boxed{\phi \;=\; 1 - R_h - \lambda_h} \qquad \Longleftrightarrow \qquad \underbrace{(1-\phi)}_{\text{trade-credit share}} \;=\; \underbrace{R_h}_{\text{cost of bank credit}} \;+\; \underbrace{\lambda_h}_{\text{shadow value of constraint}} . \tag{8}$$

Equation (8) is the central modeling object. The endogenous prepayment fraction responds to two distinct forces: the cost of bank credit R_h, and—crucially—the shadow value λ_h of the borrowing constraint. We return to this decomposition below.

Upstream sector. The upstream firm produces $Y_n = A_n L_n$ using labor L_n at wage W, and sells the intermediate at price P. To produce it must pay its wage bill $WL_n = (W/A_n)\, Y_n$ up front. It funds this from two sources: the cash prepayment $\phi\, PX$ received from the downstream firm, and bank borrowing, the latter capped at a fraction θ_n of its sales:

$$\frac{W}{A_n}\, Y_n \;\leq\; \phi\, PX \;+\; \theta_n\, P\, Y_n. \tag{9}$$

The prepayment ϕPX relaxes the upstream constraint: cash received from the buyer substitutes for the firm's own scarce borrowing. In the constrained regime—which we assume throughout, and which requires the per-unit wage cost to exceed the per-unit collateral value, $W/A_n > \theta_n P$—constraint (9) binds and pins down the upstream firm's feasible output:

$$Y_n \;=\; \frac{\phi\, PX}{W/A_n - \theta_n P}. \tag{10}$$

Upstream output is increasing in the prepayment it receives: a larger cash advance lets the supplier assemble more working capital and produce more.

4.1.2 Equilibrium and the Transmission Mechanism

A (partial) equilibrium is a pair of optimal firm choices at which both borrowing constraints (7) and (9) bind. The exogenous shock is a tightening of the downstream borrowing limit θ_h. The following proposition characterizes how this shock transmits to the upstream sector.

Proposition 1 (Endogenous trade credit and upstream transmission) *In the constrained regime, the two-sector equilibrium satisfies:*

(i) (Endogenous prepayment.) *The downstream firm's optimal prepayment is* $\phi = 1 - R_h - \lambda_h$ *with* $\lambda_h > 0$*; equivalently the trade-credit share is* $1 - \phi = R_h + \lambda_h$*.*

(ii) (Distress lowers prepayment.) *A tightening of the downstream borrowing limit raises the shadow value of the constraint and reduces prepayment:*

$$\frac{d\lambda_h}{d\theta_h} < 0, \qquad \frac{d\phi}{d\theta_h} > 0.$$

(iii) (Transmission through payment terms.) *Upstream output is strictly increasing in the prepayment it receives,* $\partial Y_n / \partial \phi > 0$*. Combining with (ii), the downstream shock transmits to lower upstream output through the prepayment channel: holding the downstream's order fixed,*

$$\left.\frac{\partial Y_n}{\partial \theta_h}\right|_{\text{prepayment channel}} = \frac{\partial Y_n}{\partial \phi}\,\frac{d\phi}{d\theta_h} > 0,$$

so that $\theta_h \downarrow \Rightarrow\ \phi \downarrow \Rightarrow\ Y_n \downarrow$*.*

(iv) (No channel under exogenous trade credit.) *If* ϕ *is held fixed at an exogenous value (the trade-credit fraction is a parameter), then* $d\phi/d\theta_h = 0$ *and the channel in (ii)–(iii) vanishes: the downstream firm's financial distress cannot reach the upstream sector through payment terms.*

The proof is in Appendix B.

Mechanism. The transmission chain is $\theta_h \downarrow \Rightarrow\ \lambda_h \uparrow \Rightarrow\ \phi \downarrow \Rightarrow\ \phi P X \downarrow \Rightarrow\ Y_n \downarrow$. A tightening of the downstream firm's borrowing limit raises the shadow cost of bank credit; by (8) the firm responds by cutting prepayment and leaning more on trade credit—exactly the payment-delay behavior

documented in Section 3, where constrained real-estate firms expanded accounts payable. Because the supplier finances production in part out of that prepayment, the reduced cash advance tightens the supplier's own constraint and forces it to curtail output. The financial shock thus travels upstream not through a fall in physical demand but through the *terms of payment*.

Contrast with exogenous trade credit. Part (iv) isolates our departure from the canonical production-network treatment of trade credit. In Altinoglu (2021), the trade-credit fraction is an exogenous parameter: the trade-credit *amount* moves mechanically with revenue, but the *terms* do not respond to financial conditions, so a borrowing-constraint shock cannot propagate through payment delays. Endogenizing ϕ through the financing problem (8) opens precisely this channel.

Contrast with cost-of-credit theories. Equation (8) also distinguishes our mechanism from the endogenous-trade-credit literature that derives payment terms from cost-of-credit minimization (Reischer, 2020; Luo, 2020; Cun et al., 2022). There, trade credit is in effect an alternative priced form of borrowing, and its use responds to the relative cost of bank versus supplier credit—the R_h term in (8). In our setting, by contrast, prepayment responds to the shadow value λ_h of the *binding quantity constraint*: when bank credit is rationed, the firm is pushed toward trade credit regardless of whether interest rates move. This distinction is not merely formal. The Three Red Lines policy was a quantity restriction—caps on leverage—rather than a change in the price of credit; the relevant margin is therefore λ_h, the constraint channel that only the present formulation contains. It is also the margin consistent with our empirical findings in Section 3, in which the transmission concentrates among financially constrained firms rather than tracking interest-rate differentials.

From the mechanism to its quantification. Proposition 1 formalizes the mechanism the evidence revealed. Parts (ii) and (iii) are the comparative static behind the payment-delay behavior documented in Section 3—payables up and bank credit down at the constrained developer, receivables up and output down at its supplier—and part (iv), in which a fixed prepayment fraction shuts the channel down, is the counterfactual we make quantitative in Section 5. What the framework cannot deliver is magnitude: confined to a single transmission link and a fixed-price environment, it

makes the mechanism transparent but understates its aggregate footprint. The multi-sector model we turn to next preserves the prepayment condition (8) sector by sector and embeds it in China's 2020 input-output network, where two forces absent here—general-equilibrium adjustment of intermediate prices, and the demand-determined contraction of real-estate output—turn the transmission of Proposition 1 into quantitatively important amplification of the aggregate shock.

4.2 Multi-Sector Model for the Chinese Economy

The two-sector model of Section 4.1 above isolates the transmission mechanism in its simplest form. To quantify the policy's aggregate effect, we now embed this mechanism in a multi-sector economy with input-output linkages, calibrated to the Chinese real-estate context.

The economy has seven sectors. Six produce commodity goods that serve as intermediate inputs to other sectors and as final consumption goods for households: Agriculture, Mining, Basic Manufacturing, Advanced Manufacturing, Construction, and Services (indexed by i or j). The seventh sector, indexed by h, is Real Estate. Real Estate plays a distinctive role in the model, dictated by three features of its output and demand. First, new housing is a durable asset that households accumulate over time, rather than a flow good consumed contemporaneously. Second, other commodity sectors do not use new housing as an intermediate input: the flows recorded from real estate to industry in the input-output table primarily reflect commercial leasing services and imputed rental income on owner-occupied housing, not productive intermediate use of new construction. Third, households' demand for housing depends on a time-varying preference weight χ_t that we model as responding endogenously to credit conditions in real estate. We detail each block of the model in turn.

4.2.1 Households

The representative household maximizes

$$\max \ \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \left[\ln C_t + \chi_t \ln H_t - \frac{L_t^{1+\eta}}{1+\eta} \right] \tag{11}$$

over aggregate consumption C_t, the housing stock H_t, and labor supply L_t. The preference weight χ_t has steady-state value $\bar{\chi}$; its dynamics are specified in Section 4.2.4. The flow budget constraint

is

$$P_t C_t + Q_t H_t = Q_t(1-\delta_h)H_{t-1} + W_t L_t + \Pi_t \tag{12}$$

where Q_t is the housing asset price, δ_h is the quarterly housing depreciation rate, and Π_t is aggregate firm profits rebated to households. The first-order conditions yield labor supply $L_t^{\eta} = W_t/(P_t C_t)$ and the housing Euler equation

$$\frac{Q_t}{P_t C_t} \;=\; \frac{\chi_t}{H_t} \;+\; \beta\,(1-\delta_h)\,\mathbb{E}_t\left[\frac{Q_{t+1}}{P_{t+1}C_{t+1}}\right], \tag{13}$$

which pins Q_t down as the marginal-utility-weighted present value of a unit of housing stock. Aggregate consumption is a Cobb-Douglas aggregator over commodity goods (real estate excluded, since households derive housing utility from the stock H_t rather than from a current flow of housing services):

$$C_t = \prod_i \left(\frac{C_{i,t}}{\gamma_i}\right)^{\gamma_i}, \qquad P_t = \prod_i P_{i,t}^{\gamma_i},$$

with $\sum_i \gamma_i = 1$. Demand for individual commodities satisfies $P_{i,t}C_{i,t} = \gamma_i P_t C_t$.

4.2.2 Commodity Sectors

Timing and information. Within each period, inputs, labor, and the prepayment fraction are committed before output is sold and financed against that period's sales receipts—the working-capital structure behind the borrowing constraints below. The only forward-looking, ex-ante uncertain object is the house price Q_t: it enters the household's housing Euler equation (13) through $\mathbb{E}_t[Q_{t+1}]$, and the real-estate developer commits production against its expected sale value $\mathbb{E}_t[Q_t]$. Commodity-sector prices, by contrast, are pinned down by within-period market clearing, so each commodity firm solves an effectively static within-period problem; the household, maximizing expected utility (11), supplies the only intertemporal link. Because the economy's only disturbance is the one-time policy innovation ε^{θ} introduced below, with no further shocks thereafter, we solve for the perfect-foresight response: along the equilibrium path every expectation operator equals its realized value, so the first-order conditions below—and the housing Euler equation (13)—are written in realized-price form.

Each commodity sector i produces using labor and a composite intermediate input:

$$Y_{i,t} = A_{i,t}\, L_{i,t}^{\alpha_i}\, X_{i,t}^{1-\alpha_i}, \qquad X_{i,t} = \prod_{j \neq i} \left(\frac{X_{ij,t}^{C}}{\zeta_{ij}} \right)^{\zeta_{ij}},$$

with $\sum_{j \neq i} \zeta_{ij} = 1$. The exclusion of real estate from the intermediate aggregator ($\zeta_{ih} = 0$) reflects the empirical observation that commodity sectors do not consume new housing as a productive input. The price index of the intermediate bundle is $S_{i,t}^{C} = \prod_{j \neq i} P_{j,t}^{\zeta_{ij}}$ and the within-sector intermediate demand satisfies $P_{j,t} X_{ij,t}^{C} = \zeta_{ij} S_{i,t}^{C} X_{i,t}$.

Sector i chooses a prepayment fraction $\phi_{i,t}$ on its intermediate purchases, paying $\phi_{i,t}$ upfront and deferring $1 - \phi_{i,t}$ as trade credit. Trade credit incurs a default cost $G(\phi_{i,t})\, S_{i,t}^{C} X_{i,t}^{C}$ with $G(\phi_{i,t}) = \frac{1}{2}(1 - \phi_{i,t})^2$. The firm's financing constraint requires that net cash outflows be financed by bank credit subject to the limit θ_i:

$$\begin{aligned} &\underbrace{\phi_{i,t} S_{i,t}^{C} X_{i,t}^{C}}_{\text{prepayment to suppliers}} + \underbrace{W_t L_{i,t}}_{\text{wages}} \\ &- \underbrace{\sum_{j \neq i} \phi_{j,t} \zeta_{ji} S_{j,t}^{C} X_{j,t}^{C}}_{\text{trade credit received from customers}} - \underbrace{\phi_t^{H} \zeta_{hi} S_t^{H} X_t^{H}}_{\text{trade credit received from RE}} \leq \theta_i\, P_{i,t}\, Y_{i,t}. \end{aligned} \tag{14}$$

The two negative terms on the left-hand side capture the cash that sector i receives upfront from its downstream customers — other commodity sectors that purchase good i, and the real-estate sector that purchases good i as a construction input. These inflows relax sector i's constraint. When a customer reduces its ϕ, the inflow shrinks and sector i's effective financing need rises — the channel through which financial pressure propagates upstream.

The firm maximizes profits subject to the production function and (14). The first-order conditions are

$$(1 + R_i + \mu_{i,t})\, W_t L_{i,t} = (1 + \theta_i \mu_{i,t})\, \alpha_i\, P_{i,t}\, Y_{i,t} \tag{15}$$

$$\left[1 + R_i \phi_{i,t} + G(\phi_{i,t}) + \mu_{i,t} \phi_{i,t}\right] S_{i,t}^{C} X_{i,t}^{C} = (1 + \theta_i \mu_{i,t})\, (1 - \alpha_i)\, P_{i,t}\, Y_{i,t} \tag{16}$$

$$\phi_{i,t} = 1 - R_i - \mu_{i,t} \tag{17}$$

where $\mu_{i,t} \geq 0$ is the Lagrange multiplier on (14), with complementary slackness $\mu_{i,t} \cdot \text{slack}_{i,t} = 0$.

Equation (17) is the multi-sector analog of (8) in the two-sector framework: when the constraint binds harder ($\mu_{i,t}$ rises), the firm reduces prepayment ($\phi_{i,t}$ falls).

4.2.3 Real Estate Sector

Real estate produces new housing IH_t using labor and a composite intermediate input drawn from the six commodity sectors:

$$IH_t = A_t^H (L_t^H)^{\alpha_h} (X_t^H)^{1-\alpha_h}, \qquad X_t^H = \prod_i \left(\frac{X_{i,t}^H}{\zeta_{hi}} \right)^{\zeta_{hi}},$$

with $\sum_i \zeta_{hi} = 1$. The output IH_t is sold exclusively to households at the housing asset price Q_t, realized at the end of the period; per the timing convention above, the developer commits production at the start of the period against the expected sale value $\mathbb{E}_t[Q_t]$, which equals Q_t under perfect foresight. New housing enters the housing stock through

$$H_t \;=\; (1-\delta_h) H_{t-1} + IH_t. \tag{18}$$

Real estate's financing block has the same structural form as that of the commodity sectors,

$$\phi_t^H S_t^H X_t^H + W_t L_t^H \;\le\; \theta_t^H \, Q_t \, IH_t. \tag{19}$$

Note that the inflow terms that appeared in (14) are absent here: real estate sells only to households (who pay cash in full upon purchase), so it does not receive trade credit from any downstream party. The first-order conditions for real estate are analogous to (15)–(17), with Q_t replacing $P_{i,t}$ on the revenue side. In particular, the trade-credit FOC retains the same structure:

$$\phi_t^H = 1 - R^H - \mu_t^H. \tag{20}$$

Real estate is a pure buyer in the trade credit network: it extends no trade credit and receives none. When θ_t^H tightens and μ_t^H rises, ϕ_t^H falls, which transmits the financial pressure to each of real estate's upstream suppliers through the inflow term $\phi_t^H \zeta_{hi} S_t^H X_t^H$ in (14). This is the trade-credit channel at the center of the paper.

4.2.4 Endogenous Housing Demand Wedge χ_t

The framework of this section is deliberately generic. Its production, financing, and trade-credit blocks contain nothing specific to housing—which is what allows the same machinery to be carried to bank deleveraging, to sectoral environmental regulation, or to distress at a firm near the center of a supply network, and portability is part of what we claim for it. Genericity has a price at the point of application. A production-network model with durable housing can represent the Three Red Lines in one way only: as a restriction on the supply of new housing. Read that way, the model's price prediction is not simply wrong but informative. Under a constant preference weight, a tightening of θ_t^H lowers IH_t, pushes H_t below trend, raises the term χ_t / H_t, and with it Q_t—the standard consequence of restricting the supply of a durable good. China's 70-city new-home price index instead declined by approximately 2.3% over the eight quarters following 2020Q3. The discrepancy locates what the production side omits.

What it omits is not a taste for housing. It is that Chinese households do not buy houses; they buy claims on houses not yet built. As Section 2.3 notes, pre-sale deposits and the mortgages that fund them supplied roughly half of developers' funds in place in 2020, making the household at once the sector's largest financier and the party that bears delivery risk. What the household purchases is therefore a claim whose value depends on the developer's solvency, and a regulation that impairs that solvency degrades the asset being sold: a completed apartment and a promise of one from a distressed developer are different goods, and they command different prices. Households repriced accordingly. By mid-July 2022, buyers at more than one hundred stalled projects across upwards of fifty cities had suspended mortgage payments; within five weeks the count exceeded three hundred projects spanning twenty-four of China's thirty-one provinces, and the authorities' "ensure delivery" campaign followed as the official response.[7]

χ_t is therefore the application-specific layer of an otherwise industry-neutral framework: the one point at which the institutional features of Chinese housing enter. We summarize the dependence by letting χ_t evolve as an AR(1) driven by the credit shock:

$$\log\left(\frac{\chi_t}{\bar{\chi}}\right) \;=\; \rho_\chi \log\left(\frac{\chi_{t-1}}{\bar{\chi}}\right) \;+\; (1-\rho_\chi)\,\kappa_\chi \,\log\left(\frac{\theta_t^H}{\bar{\theta}^H}\right), \tag{21}$$

[7] China Real Estate Information Corporation, analysis of 13 July 2022, and contemporaneous press accounts for the August figures.

where $\kappa_\chi > 0$ controls the long-run elasticity of χ_t with respect to the shock to real-estate borrowing constraints and $\rho_\chi \in [0, 1)$ controls persistence. Alongside delivery risk, (21) absorbs two further forces: tighter developer credit thins the resale market for housing households already own, and the policy signaled a lasting shift in the government's stance toward real estate, weakening long-run demand—persistently, though $\rho_\chi < 1$ implies not permanently. None of the three is micro-founded here; χ_t stands in for their net effect on household demand.

Three features keep the specification disciplined. First, χ_t carries no innovation of its own. It is a deterministic function of θ_t^H, so the model retains a single exogenous driving force. Set against the housing preference shocks of Iacoviello (2005) and Liu et al. (2013), which are independent stochastic processes, the wedge is more constrained rather than less: it cannot move except in response to the policy. Second, its two parameters are pinned by two house-price moments—the cumulative eight-quarter decline and its persistence—neither of which is a quantity this paper claims. Third, a micro-foundation of the delivery-risk channel, built on pre-sale contracts, beliefs about developer default, and the mortgage claim, would deliver an object of precisely this form: a wedge in the housing Euler equation that co-moves with developer distress. We write that wedge down directly and label it for what it is. Supplying the micro-foundation would sharpen its interpretation, not its mechanics, and we leave that to future work.

None of this, finally, is load-bearing for the paper's central claim. χ_t enters the model only through the household's housing Euler equation (13). It does not appear in any production-side equation, firm financing constraint, or trade-credit first-order condition. That it touches exactly one equation is the boundary between the portable core and the application-specific layer, and the placement is deliberate: a reader carrying this framework to another sector discards (21) and changes nothing else. The trade-credit mechanism therefore operates independently of χ_t, and the counterfactual exercises in Section 5 can shut down the trade credit channel while leaving χ_t active. In particular, the steady-state trade-credit block—the multipliers μ_i and the prepayment fractions ϕ_i—is invariant to the level of $\bar{\chi}$, and the counterfactual decomposition in Section 5 holds the path of χ_t fixed across scenarios, so the roughly 40% share that the decomposition attributes to trade credit is unaffected by the wedge.

4.2.5 Equilibrium

A perfect-foresight equilibrium is a sequence of prices $\{P_{i,t}, Q_t, W_t\}$, real allocations $\{C_{i,t}, X^C_{ij,t}, X^H_{i,t}, L_{i,t}, L^H_t, IH_t, H_t\}$, financial variables $\{\mu_{i,t}, \mu^H_t, \phi_{i,t}, \phi^H_t\}$, and the wedge χ_t, given the exogenous shock path $\{\theta^H_t\}$ and the predetermined housing stock H_{t-1}, such that: households optimize; commodity sectors and real estate optimize; complementary slackness holds for each borrowing constraint; χ_t follows (21); commodity goods markets clear for $i \neq h$,

$$Y_{i,t} = C_{i,t} + X^H_{i,t} + \sum_{j \neq i} X^C_{ji,t};$$

the housing market clears through (18); the labor market clears, $L_t = L^H_t + \sum_i L_{i,t}$; and a price normalization closes the nominal scale. The solution method uses a Fischer-Burmeister formulation of the complementary slackness conditions and a forward–backward iteration on Q_t. The model is thus *semi-dynamic*: given the path of Q_t, each period is a static equilibrium, and it is the motion of Q_t—the only forward-looking price—that moves the prepayment fractions $\phi_{i,t}$ and ϕ^H_t, and hence trade credit, over time.

4.2.6 Steady State

Before introducing the policy shock, consider the economy with no policy change: $\theta^H_t = \bar{\theta}^H$ and $\chi_t = \bar{\chi}$ for all t. The equilibrium is then stationary—all prices, allocations, and financial variables are constant—and two properties of the objects at the center of the analysis follow immediately.

First, the *house price is constant*. Imposing $Q_t = Q_{t+1} = Q$ and constant consumption on the housing Euler equation (13) gives the closed form

$$\frac{Q}{PC} = \frac{\bar{\chi}/H}{1 - \beta(1 - \delta_h)}, \tag{22}$$

so Q is pinned down by the preference weight $\bar{\chi}$, the housing stock H, and discounting. The stock is itself stationary, which by (18) requires replacement investment $IH = \delta_h H$.

Second, *trade credit is constant*: each sector's prepayment fraction is $\phi_i = 1 - R_i - \mu_i$ (and $\phi^H = 1 - R^H - \mu^H$), with $\mu_i, \mu^H \geq 0$ the steady-state shadow values of the borrowing limits. Real estate is a *pure buyer*—it receives no trade-credit inflow to relax its constraint, and finances against

its (relatively small) housing collateral QIH—so it carries the highest multiplier μ^H, and hence the lowest prepayment ϕ^H, already at the steady state.

In the calibrated network it is natural to read ϕ_i as a *net cash-payment position* rather than the bounded fraction $[0, 1]$ of the two-sector framework in Section 4.1. Numerically, $Q \approx 0.43$ and the commodity prepayments lie in $[0.26, 0.95]$, while real estate's ϕ^H is the lowest—and in fact slightly *negative*. A negative ϕ^H says that the dominant, severely constrained buyer *net-receives* short-term financing from its suppliers; this is not an artifact but the model counterpart of our central empirical fact—rising payables at constrained developers mirrored by rising receivables at their suppliers, the pre-sale-funded developer running on supplier credit.

We solve the steady state numerically by random-restart `fsolve`; it is the $t \to \infty$ anchor for the perfect-foresight transition that the policy shock sets off.

4.3 Calibration

We now bring the model to the data. We calibrate it at quarterly frequency to the Chinese economy on the eve of the Three Red Lines policy (2018Q1–2020Q2). The calibration draws on three sources: China's 2020 national input-output table at the 153-sector level, aggregated to our seven sectors; firm-level financial statements from CSMAR; and standard parameters from the macro-finance literature. We organize the discussion into four blocks—preferences and depreciation, the structure of the production network, financial parameters, and the shock process and demand wedge—and summarize all parameter values in Tables 8 and 9.

4.3.1 Preferences and Depreciation

The discount factor $\beta = 0.99$ corresponds to a quarterly model with an annual real interest rate of approximately 4%. The inverse Frisch elasticity is $\eta = 0.5$, in the middle of the range standard in housing-cycle DSGE models. The quarterly housing depreciation rate is $\delta_h = 0.01$, consistent with the annual depreciation of approximately 4% that Iacoviello and Neri (2010) use for residential housing. The intertemporal elasticity of substitution is unity ($\sigma = 1$, log utility over C). We choose the steady-state housing preference weight $\bar{\chi} = 0.40$ so that the steady-state ratio of housing wealth to consumption, $QH/(PC)$, is in line with the Chinese household balance-sheet aggregates

documented in Liu et al. (2013).

4.3.2 Production-Network Structure

We aggregate the 2020 input-output table from 153 sectors into seven sectors aligned with the CSMAR industry classification used in the empirical analysis: Agriculture (A), Mining (B), Basic Manufacturing (C13–C31), Advanced Manufacturing (C32–C43), Construction (E), Real Estate (K70), and Services (D, F–I, L–R). The full aggregation map is reported in Appendix C.

Labor shares. Each sector's labor share α_i is computed from the input-output table as labor compensation divided by total cost (labor plus intermediate inputs). Real estate and agriculture have the highest labor shares ($\alpha_h = 0.72$, $\alpha_{\text{Agri}} = 0.62$), reflecting their labor-intensive production. Manufacturing sectors have the lowest labor shares (0.19 for Advanced Manufacturing and 0.25 for Basic Manufacturing), reflecting their high intensity of intermediate-input use.

Consumption shares. Final consumption shares γ_i are computed from the input-output table's household final-demand column, normalized over commodity sectors so that $\sum_i \gamma_i = 1$. The consumption basket is dominated by Services ($\gamma_{\text{Serv}} = 0.69$) and Basic Manufacturing (0.19), with much smaller shares for Agriculture (0.07), Advanced Manufacturing (0.06), and trace shares for Mining and Construction. Real estate's consumption share is set to zero by construction — housing services are obtained from the durable stock H_t rather than from contemporaneous flow purchases.

Intermediate input shares. The intermediate input share matrix ζ_{ij} is constructed from the input-output table's coefficients of intermediate use. We then set the real-estate column to zero ($\zeta_{ih} = 0$ for all i), reflecting the model's treatment of new housing as a final-use good, and renormalize each row to sum to one over the remaining six commodity sectors. Real estate's own input share vector ζ_{hi} is taken directly from the input-output table without modification: it shows that real estate's largest input requirements come from Services (0.74), Basic Manufacturing (0.06), and Construction (0.05). The complete ζ matrix is reported in Appendix D.

4.3.3 Borrowing Constraints and Interest Rates

The steady-state borrowing limit $\bar{\theta}_i$ governs the share of revenue that can be financed via bank credit and is the key financial parameter at the sector level. We calibrate $\bar{\theta}_i$ from CSMAR firm-level data on the median ratio of bank loans (short-term plus long-term) to sales revenue over the pre-policy window (2018Q1–2020Q2). The sector medians range from 0.26 (Basic Manufacturing) to 1.68 (Real Estate), confirming real estate's distinctive reliance on bank financing and consistent with the policy concern that motivated the Three Red Lines intervention.

Two considerations shape the mapping from raw data to model parameters. First, the model parameter $\bar{\theta}_i$ should be such that the borrowing constraint binds at steady state with positive Lagrange multiplier $\mu_i > 0$; this is required for the trade-credit mechanism (recall $\phi_i = 1 - R_i - \mu_i$) to operate meaningfully. Second, the cross-sector ordering observed in the data — with real estate most constrained and basic manufacturing least — is a substantive feature we want to preserve. We accommodate both by linearly mapping the raw data into the interval $[\theta_{\min}, \theta_{\max}] = [0.30, 0.50]$:

$$\bar{\theta}_i \;=\; \theta_{\min} + \frac{\theta_i^{\text{data}} - \min_j \theta_j^{\text{data}}}{\max_j \theta_j^{\text{data}} - \min_j \theta_j^{\text{data}}} \cdot (\theta_{\max} - \theta_{\min}).$$

The resulting $\bar{\theta}_i$ values are reported in Table 9. Real estate attains the upper bound ($\bar{\theta}_h = 0.50$) and Basic Manufacturing the lower bound (0.30); the remaining sectors lie in between. The interval $[0.30, 0.50]$ is chosen so that all steady-state constraint multipliers μ_i are strictly positive but moderate, which keeps the model's impulse responses well behaved.

The bank interest rate is $R_i = 0.05$ across sectors and time. We interpret this as a quarterly financing premium relative to the household discount rate; the parameter affects the level of steady-state trade credit ($\bar{\phi}_i = 1 - R_i - \bar{\mu}_i$) but has limited influence on the dynamics of the impulse response, which depend on deviations from steady state.

4.3.4 Shock Process and Demand Wedge

Real-estate credit shock. The real-estate borrowing limit follows a stochastic AR(1) process in logs,

$$\log \theta_t^H = (1-\rho) \log \bar{\theta}^H + \rho \, \log \theta_{t-1}^H + \varepsilon_t^\theta, \qquad \rho = 0.8,$$

where ε_t^θ is a mean-zero innovation to the regulatory stance: agents do not perfectly foresee future deleveraging intensity, so policy is a genuine source of uncertainty. We model the Three Red Lines as a single, unanticipated realization $\varepsilon_1^\theta = -\sigma_{\text{shock}}$ (with $\sigma_{\text{shock}} = 0.3$ and $\varepsilon_t^\theta = 0$ for $t \neq 1$) and trace the economy's perfect-foresight response to it—an "MIT shock" that tightens the limit at $t = 1$ and reverts at rate ρ. The peak deviation is $\log(\theta_1^H/\bar{\theta}^H) = -0.30$, corresponding to a 25.9% decline in the borrowing limit. We calibrate σ_{shock} to match the cumulative contraction in real estate developers' ratio of bank loans to revenue between 2019 and the trough of the credit decline in 2021–2022, as documented in the aggregate CSMAR data (Section 3). The persistence $\rho = 0.8$ matches the half-life of the credit decline observed in the same data (approximately 3 quarters to reach half the peak deviation).

Housing demand wedge. The wedge χ_t follows the AR(1) in (21). We calibrate $(\kappa_\chi, \rho_\chi) = (1.5, 0.92)$ to match two features of China's 70-city new-home price index over 2020Q3–2022Q4: (i) the eight-quarter cumulative price decline of approximately 2.3%, and (ii) the persistence of the decline over the subsequent two-year horizon. The value $\rho_\chi = 0.92$ is at the high end of the range used in the literature on housing preference shocks (Iacoviello and Neri, 2010; Liu et al., 2013), consistent with the slow recovery of Chinese housing prices in our sample.

4.3.5 Parameter Summary

Table 8 reports all non-sectoral parameters and Table 9 reports sector-specific values for α_i, γ_i, and $\bar{\theta}_i$. The complete ζ_{ij} matrix and the mapping from the input-output table to the seven sectors are deferred to Appendix D. One feature of the calibration deserves emphasis: no moment we target is a spillover moment. The network structure comes from the input-output table, the financial block from pre-policy balance sheets, the shock from developers' observed credit contraction, and the wedge from house prices. The cross-sector transmission the model generates is therefore a prediction, not a target.

4.4 Baseline Impulse Response

Figure 2 reports the baseline impulse response of the calibrated model to the Three Red Lines shock. The six panels display percentage deviations of θ_t^H, the housing asset price Q_t, the housing

Table 8: Main calibration parameters (non-sectoral)

Parameter	Symbol	Value	Source / target
A. Preferences and depreciation			
Discount factor	β	0.99	Quarterly standard
Inverse Frisch elasticity	η	0.50	Mid-range, housing-DSGE literature
Intertemporal elasticity of subst.	σ	1.00	Log utility over C
Housing depreciation (quarterly)	δ_h	0.01	$\approx 4\%$/yr; Iacoviello and Neri (2010)
SS housing preference	$\bar{\chi}$	0.40	Targets SS $QH/(PC)$ ratio
B. Financial frictions			
Bank credit premium (quarterly)	R	0.05	Calibration of SS $\bar{\phi}$ levels
C. Shock process and demand wedge			
θ^H shock persistence	ρ	0.80	CSMAR RE bank-credit recovery profile
θ^H peak shock magnitude	σ_{shock}	0.30	-25.9% peak θ^H deviation
χ elasticity to θ^H	κ_χ	1.50	8q cumulative Q decline $= -2.3\%$
χ persistence	ρ_χ	0.92	70-city price persistence, 2020Q3–2022Q4

Notes: Non-sectoral parameters of the multi-sector model. Sector-specific values (labor shares α_i, consumption shares γ_i, steady-state borrowing limits $\bar{\theta}_i$) are reported in Table 9. The intermediate input share matrix ζ_{ij} is constructed from the 2020 China IO table with $\zeta_{i,h} = 0$ imposed (commodity sectors do not use real estate as an intermediate input); the full 7×7 matrix is reported in Appendix D.

stock H_t, new housing investment IH_t, aggregate consumption C_t, and aggregate GDP from their respective steady-state values, over the first 120 quarters after the shock.

Four features of the baseline response are worth highlighting. First, Q_t declines on impact and reaches a cumulative eight-quarter decline of -2.3%, by calibration design. Second, new housing investment IH_t falls sharply (trough $\approx -25\%$ at $t = 1$, gradual recovery), reflecting both real estate's tighter financing constraint and weaker housing demand. Third, the housing stock H_t accumulates the flow shortfall and remains persistently below trend for many years, consistent with the durable-asset structure. Fourth, aggregate GDP declines by approximately 4.7% at trough and recovers gradually, for a cumulative 16-quarter loss of -30.9 percentage-point-quarters. The counterfactual decomposition in Section 5, which builds on this impulse response, reports the cumulative loss over the 20-quarter window that matches the regression sample.

Figure 3 reports the impulse response of the two financial variables central to the trade-credit mechanism: the borrowing-constraint multiplier μ_t^H and the prepayment fraction ϕ_t^H for real estate. The shock causes μ_t^H to rise sharply on impact (the constraint binds harder), and through the first-order condition (20) the prepayment fraction ϕ_t^H falls by a corresponding amount. The decline in ϕ_t^H is the upstream-facing object that operationalizes the trade-credit channel: when

Table 9: Sector-specific calibration parameters

Sector	α_i (labor share)	γ_i (cons. share)	θ_i^{data} (CSMAR median)	$\bar{\theta}_i$ (mapped)
Agriculture	0.617	0.070	0.536	0.339
Mining	0.540	0.000	0.516	0.336
Basic Manufacturing	0.250	0.187	0.260	0.300
Adv. Manufacturing	0.194	0.055	0.286	0.304
Construction	0.252	0.000	0.550	0.341
Real Estate	0.722	—	1.684	0.500
Services	0.490	0.688	0.325	0.309

Notes. α_i (labor share) is computed from the 2020 IO table as labor compensation divided by total cost (labor plus intermediates). γ_i (consumption share) is computed from the IO table's household final-demand column, normalized over commodity sectors so that $\sum_i \gamma_i = 1$; real estate's consumption share is zero by construction because housing services accrue from the durable stock H_t rather than from a current flow. θ_i^{data} is the sector-median ratio of bank loans (short-term plus long-term) to sales revenue in CSMAR, computed over the pre-policy window 2018Q1–2020Q2. $\bar{\theta}_i$ is the model-implied steady-state borrowing limit, obtained by linearly mapping θ_i^{data} into the interval $[\theta_{\min}, \theta_{\max}] = [0.30, 0.50]$. The mapping preserves cross-sector ordering and yields strictly positive SS constraint multipliers $\bar{\mu}_i > 0$ for all sectors, which is required for the endogenous trade-credit mechanism ($\phi_i = 1 - R - \mu_i$) to operate meaningfully.

real estate reduces ϕ_t^H, its suppliers receive a smaller fraction of their invoices upfront and absorb the shortfall on their own balance sheets. The magnitude and persistence of the ϕ_t^H decline therefore determine the quantitative importance of the trade-credit channel in the aggregate response. The decomposition in Section 5 shows that this channel accounts for roughly 40% of the total cumulative GDP loss generated by the shock.

5 Quantitative Analysis

This section puts the calibrated model to work. We first present the paper's headline quantitative result: a counterfactual decomposition that attributes roughly 40% of the aggregate response to the endogenous adjustment of trade credit (Section 5.1), followed by an over-identification check that confronts the model with the reduced-form estimates. We then ask where the loss lands across sectors (Section 5.2), evaluate three policy counterfactuals (Section 5.3), and assess the robustness of the quantitative findings (Section 5.4).

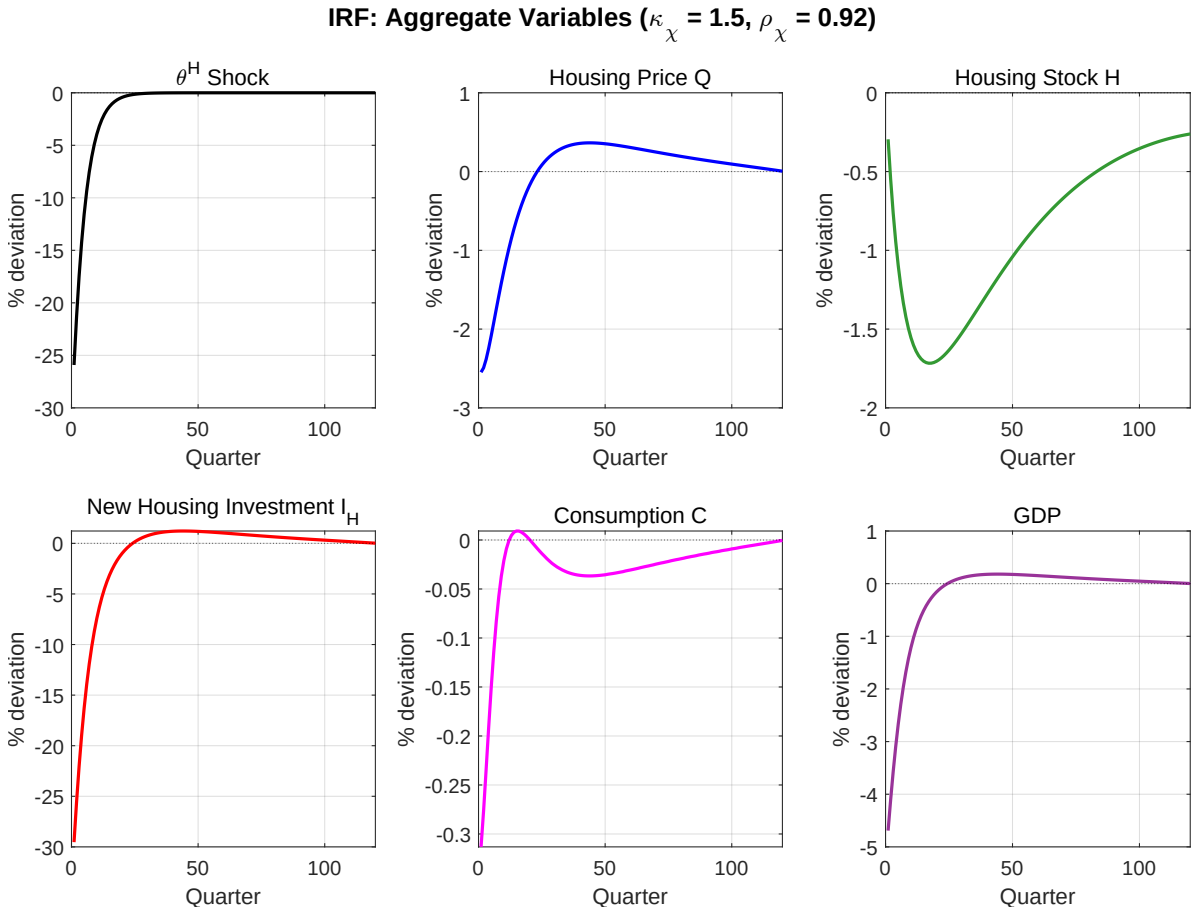


Figure 2

Notes: Baseline impulse response to a one-time tightening of the real estate borrowing constraint θ_t^H. Panels show percentage deviations from steady state for the shock itself, the housing asset price Q_t, the housing stock H_t, new housing investment IH_t, aggregate consumption C_t, and aggregate GDP. The shock is calibrated to peak at -25.9% and decays with quarterly AR(1) persistence $\rho = 0.8$; the induced path of the housing demand wedge χ_t follows (21) with $(\kappa_\chi, \rho_\chi) = (1.5, 0.92)$.

5.1 Quantifying the Channel: Decomposition and an Over-Identification Check

The decomposition. How much of the aggregate loss does the endogenous response of trade credit carry? To answer this question, we feed the same shock through two versions of the model. In Scenario A—the baseline—the prepayment fraction ϕ adjusts endogenously in every sector through the first-order condition that ties it to the shadow value of the borrowing constraint. In Scenario B, we hold ϕ at its steady-state value in every sector and period, severing it from the borrowing constraint; Scenario B is the model a researcher would obtain by treating trade credit as exogenous, as in Altinoglu (2021). Everything else—preferences, technology, the production

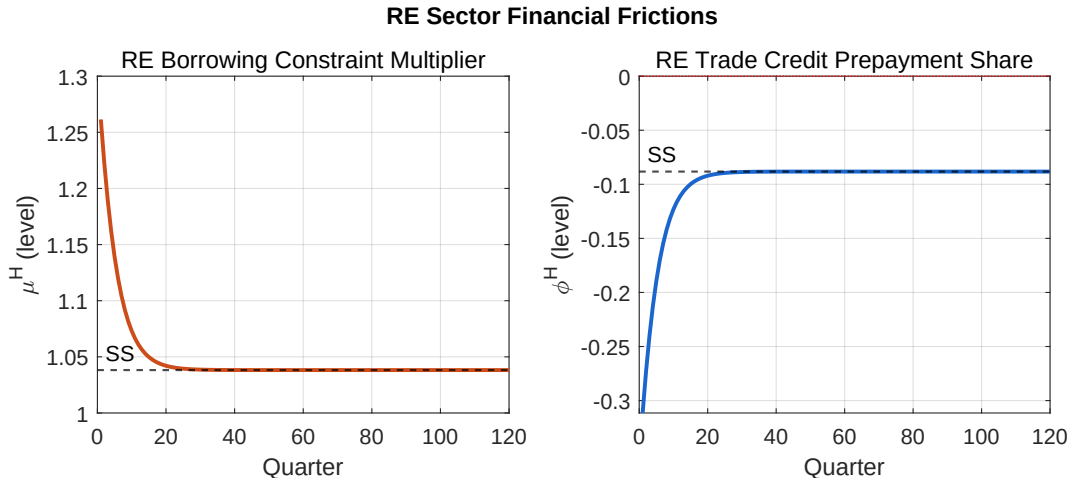


Figure 3

Notes: Baseline impulse response of real estate's financial variables. *Left panel:* the borrowing-constraint multiplier μ_t^H, which rises on impact and decays slowly back to its steady-state value (dashed line). *Right panel:* the prepayment fraction ϕ_t^H, which falls by the same amount as μ_t^H rises through the firm's first-order condition $\phi_t^H = 1 - R^H - \mu_t^H$. The decline in ϕ_t^H operationalizes the trade-credit channel by transferring liquidity pressure from real estate to its upstream suppliers.

network, the path of θ_t^H, and the induced path of the housing demand wedge χ_t—is identical across the two scenarios, so the gap between them is, by construction, the contribution of the endogenous response of trade credit. The construction also makes the decomposition agnostic about the wedge's deeper origin: whether the housing-demand decline ultimately reflects the policy itself, delivery risk, or a longer-run revaluation of housing demand, it enters both scenarios identically and nets out of the comparison.

The gap is large. Over the 20-quarter window that matches our regression sample, the cumulative GDP loss is -31.9 pp·q in Scenario A and -19.1 pp·q in Scenario B. The endogenous adjustment of trade credit thus accounts for about 40% of the aggregate response—a $1.7\times$ amplification. This decomposition is invariant to the comparison window and to the model's shock persistence, and it does not rely on the precise magnitude of the empirical estimates. It is the result we emphasize.

Why compare the model to the reduced form at all? The decomposition is internal to the model. A natural question is whether the model's aggregate response is even in the right range—and here the reduced-form estimates of Section 3 provide an external benchmark, though they speak

in different units: the empirical coefficients are in 100M per firm per intensity percentage-point, while the model produces percentage deviations of aggregate GDP from its pre-shock steady state. Before bridging the units, we state why the comparison is informative. The argument has three steps.

First, the two objects are *comparable* because they ride the same edges. Trade credit exists only where a real trade relationship exists, so dense linkage through trade credit *is* dense real linkage. Trade credit is the financial friction on the production network's own transaction edges, not a parallel credit channel: the network supplies the edges, and the friction carries the financial shock. The transmission that the regressions pick up and the transmission that the model computes are therefore the same economic object, viewed at two resolutions.

Second, the two measurements are *independent*. The reduced form cleanly identifies the channel's existence, sign, and incidence, but it cannot recover the general-equilibrium magnitude: the industry-by-time fixed effects that secure identification also absorb the aggregate adjustment, so the estimate is partial-equilibrium by construction—this is precisely why we build the model. The model, in turn, is calibrated to moments that the regressions never touch: the 2020 input–output table, sector-level borrowing limits, and the observed contraction in bank credit to developers. No spillover moment is targeted.

Third, if two routes built on disjoint identifying information arrive at losses of the same order of magnitude, the agreement serves as an over-identifying check on the mechanism. The conclusion then rests on neither route alone—and the headline decomposition above does not depend on this check at all.

Construction of the EILR. Let β_S denote the "Sales" coefficient in column (2) of Panel A of Table 2, expressed in 100M per firm-quarter per intensity percentage-point. Let $\bar{I}$ be the regression sample mean of "Intensity" (in pp), $\bar{S}$ the post-period average quarterly sales per listed firm (in 100M), and λ the share of nominal GDP accounted for by listed-firm sales. Aggregating the per-firm marginal effect across firms and scaling to economy-wide GDP yields a per-quarter implied GDP loss, the Empirical Implied Loss Ratio (EILR):

$$\mathrm{EILR}_q \;=\; \frac{\beta_S \cdot \bar{I} \cdot \lambda}{\bar{S}} \times 100, \qquad \mathrm{EILR}_{\mathrm{cum}} \;=\; Q_{\mathrm{post}} \cdot \mathrm{EILR}_q, \tag{23}$$

Table 10: Reconciling Model GDP Loss with Empirical Magnitudes

	Per-quarter GDP loss (%)	Cumulative 20q (pp·q)	Share of EILR (%)
Panel A. Empirical implied magnitude			
EILR (Sales)	−2.02	−40.43	100.0
Panel B. Model predictions (20-quarter horizon)			
(A) Baseline (ϕ endogenous)	−1.59	−31.87	78.8
(B) Trade credit off ($\phi = \bar{\phi}$)	−0.96	−19.10	47.3
Trade-credit channel (A−B)	−0.64	−12.77	40.1
amplification (A/B)			1.7×

Notes. EILR (Empirical Implied Loss Ratio) aggregates the Table 2 Sales marginal effect to economy-wide GDP loss via $\text{EILR}_q = \beta_S \cdot \bar{I} \cdot \lambda / \bar{S}$, with $\beta_S = -2.04$ (100M / quarter per intensity percentage-point), $\bar{I} = 0.6925$ pp, $\bar{S} = 33.10$ 100M, and $\lambda = 0.4735$ (listed-firm sales to nominal GDP, 2021–2024 average). Cumulative figures sum percentage deviations over the 20-quarter post-policy window (2020Q4–2025Q3); the model figures truncate the IRF to the same window. "Share of EILR" is the cumulative figure divided by −40.43; the trade-credit channel row reports $(A - B)$, whose cumulative loss is 40.1% of baseline (A), an amplification of 1.7×.

where $Q_{\text{post}} = 20$ matches the post-policy quarters in our regression sample, 2020Q4 through 2025Q3. The derivation in Appendix E shows that the number of firms N cancels, leaving (23) a function of only four sample moments and one regression coefficient.

Inputs. We obtain $\beta_S = -2.04$ (column (2) of Panel A of Table 2), $\bar{I} = 0.6925$ pp, and $\bar{S} = 33.10$ 100M directly from the regression sample of Section 3. The scaling λ is computed as the ratio of total annual sales of all listed firms in the regression sample to nominal GDP, averaged over 2021–2024: listed-firm sales of 2,391.9 trillion against cumulative GDP of 5,051.2 trillion yields $\lambda = 0.4735$. Substituting into (23) gives $\text{EILR}_q = -2.02\%$ per quarter and $\text{EILR}_{\text{cum}} = -40.43$ percentage-point-quarters over the 20-quarter window.

Comparison to the model. Table 10 places the EILR next to the model's 20-quarter cumulative GDP deviation under the two scenarios. The empirically implied loss is −40.4 pp·q. Scenario A's −31.9 pp·q is of the same order of magnitude—about four-fifths over this particular window, a fraction we flag immediately as window-sensitive and do not read as a match—whereas Scenario B, the model without the endogenous channel, delivers about half. Because β_S enters nowhere in the calibration of θ_h, ν_i, or the production network, the proximity under Scenario A is an out-of-sample property: a model pinned down by pre-policy sector-level moments delivers an

aggregate response of the same order of magnitude as the reduced-form estimates imply, and the trade-credit channel is what carries it there.

Caveats. The EILR is a transparent bridge between units, not a structural estimate, and three features should be kept in mind. (i) The EILR treats the Sales coefficient as a constant per-quarter effect, whereas the model's GDP response is front-loaded and recovers over the horizon while the empirical effect is roughly flat over the post-policy window. The cumulative comparison is therefore sensitive to the window and to the model's calibrated shock persistence, even though the channel decomposition (the 40% share and the $1.7\times$ amplification) is not; we accordingly read the comparison as a same-order-of-magnitude check and headline the window-invariant decomposition. (ii) The scaling by λ assumes that listed-firm responses are representative of the economy, which may overstate the aggregate response if unlisted firms are less exposed, or understate it if smaller unlisted suppliers are more credit-constrained. (iii) We compare sums of percentage deviations rather than ratios of cumulative levels, exact only to first order; the two are numerically close here (-2.02% versus -2.31% on the levels-based denominator). One further choice works in the conservative direction: the EILR uses the Sales coefficient from Panel A of Table 2, estimated without the size control (-2.04); the size-controlled coefficient in Panel B is -3.42, which would imply a larger empirical loss and a wider gap for any model to fill. None of these caveats touches the central message: within this transparent benchmark, the endogenous response of trade credit raises the model's predicted loss by roughly 70% and moves the model from about half of the empirically implied magnitude into the same range as the empirical benchmark.

Taken together: the decomposition attributes roughly 40% of the aggregate response to the endogenous adjustment of trade credit, and an aggregation of the reduced-form estimates—built on identifying information that the calibration never sees—lands in the same order of magnitude as the model that includes the channel, and well above the model that excludes it.

5.2 Where the Loss Lands: A Sectoral Decomposition of the Trade-Credit Channel

The aggregate result of the previous subsection masks a striking heterogeneity across sectors: the same trade-credit channel that accounts for 40.9% of the total GDP loss accounts for very different shares of each individual sector's output decline. To document this, we re-run the two counter-

Table 11: Sectoral Decomposition of Output Loss into the Trade Credit Channel

Sector	Real output (%·q, 16q)			Trough	GDP contribution (pp·q)	
	Baseline cum^i_A	TC off cum^i_B	TC share (ith sector)	(%) trough^i_A	Baseline cum^i_A	TC contrib.
Real Estate	−190.25	−164.11	13.7%	−29.53	−18.32	−2.26
Construction	−106.49	−25.17	76.4%	−15.92	−0.32	−0.23
Adv. Manuf.	−17.44	−2.18	87.5%	−2.83	−0.91	−0.79
Basic Manuf.	−15.68	−2.06	86.9%	−2.52	−2.81	−2.45
Services	−13.41	−2.74	79.6%	−2.07	−7.66	−6.12
Mining	−12.14	−2.26	81.4%	−1.92	−0.40	−0.35
Agriculture	−9.52	−1.25	86.9%	−1.54	−0.46	−0.45
Aggregate	—	—	40.9%	—	−30.88	−12.64

Notes. Columns 1–2 report the sum of each sector's quarterly percentage deviation from its own steady-state output over the 16-quarter post-shock window, in two scenarios: (A) baseline with ϕ endogenous; (B) trade credit channel shut down by holding $\phi = \bar{\phi}$ at steady state. Column 3 reports the share of each sector's own output loss attributable to the trade credit channel, $(\text{cum}^i_A - \text{cum}^i_B)/\text{cum}^i_A$. Column 4 reports the deepest quarterly decline observed within the 16-quarter window. Columns 5–6 weight sector-level deviations by steady-state sales shares to yield each sector's nominal contribution to aggregate GDP deviation; the column 5 entries sum to the aggregate cumulative GDP loss of −30.88 pp·q (Scenario A), and the column 6 entries sum to the trade-credit-channel contribution of −12.64 pp·q (40.9% of total).

factual scenarios of Section 5.1 and tabulate sector-level output deviations, both as percentages of each sector's own steady-state output and as contributions to aggregate GDP. Table 11 and Figure 4 summarize the results.

The trade-credit channel dominates the spillover to upstream sectors. Across the six non-RE sectors, the trade-credit channel accounts for between 76.4% and 87.5% of the cumulative decline in real output. The manufacturing and primary-input sectors are the most TC-dependent (87.5% for advanced manufacturing, 86.9% for basic manufacturing and agriculture, 81.4% for mining), while services come slightly lower at 79.6% and construction is lowest at 76.4%. Comparing cum^i_B across these sectors shows that, absent the TC channel, upstream sectoral losses are uniformly small (between −1.3 and −2.7 sector %·q over 16 quarters, less than one-fifth of a percent per quarter on average). By itself, the Leontief production network transmits very little of the regulatory shock past real estate's boundary: the network supplies the transmission paths, but it is the trade-credit friction operating on those paths that does the work of propagation. The upstream spillover that the empirical section identifies is therefore, in the model's accounting, a financial phenomenon rather than a physical input-output phenomenon.

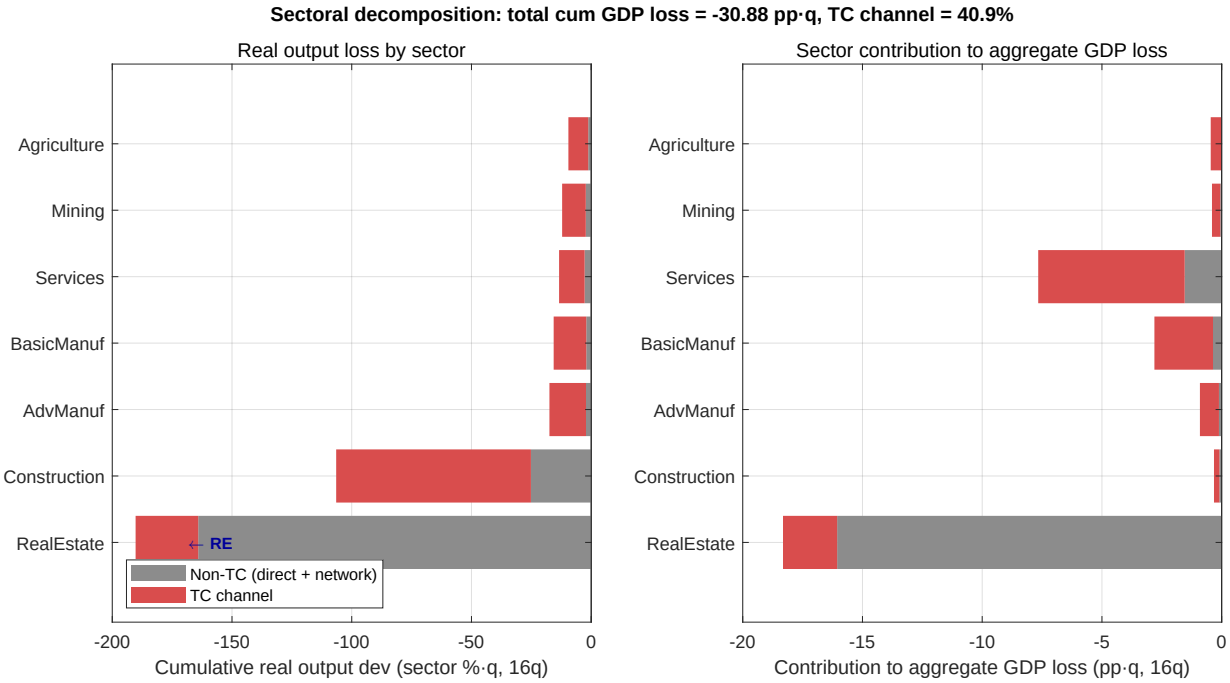


Figure 4

Notes: Sectoral decomposition of cumulative output loss (16-quarter window). *Left panel:* cumulative real output deviation for each of the seven sectors, in units of sector percent-quarter, stacked into trade credit (red) and non-TC (gray) components. *Right panel:* the same decomposition expressed as each sector's nominal contribution to aggregate GDP deviation, in percentage-points-of-GDP-quarter. Real estate is labeled in the left panel as the directly shocked sector.

Real estate displays the opposite pattern. The TC share for real estate itself is only 13.7%, the lowest of any sector. This asymmetry is a direct consequence of the model's structure: real estate is the source of the shock, not its recipient. Real estate's output loss is dominated by the two "direct" forces acting on it—the tightening of θ_h through the firm's borrowing constraint, and the contemporaneous fall in housing demand through χ_t and the asset price Q_t—both of which would operate even if ϕ were exogenous. The 13.7% TC share captures only the feedback through which upstream cuts amplify back onto real estate's own input costs and production, and is correspondingly modest.

Sectoral rank reorders sharply between intensity and aggregate weight. Two rankings emerge. Ordered by depth of own-sector impact, real estate (trough$_A = -29.5\%$) is followed by construction (-15.9%); the remaining sectors all decline by between -1.5% and -2.8% at trough. Ordered by contribution to aggregate GDP, however, services moves from a middling position to second

place (-7.66 pp·q, 24.8% of the total loss), behind only real estate (-18.32 pp·q, 59.3%) and ahead of basic manufacturing (-2.81 pp·q, 9.1%). The reordering is driven entirely by sector size: the services sector experiences only a mild per-unit decline, but its large output share magnifies its weight in any value-weighted aggregate.

Half of the trade-credit channel works through services. A corollary of the previous observation is that the aggregate footprint of the trade-credit channel is dominated by services. Services alone contributes -6.12 pp·q to the TC-channel impact—48.4% of the total TC contribution of -12.64 pp·q. Basic manufacturing contributes a further 19.4%, real estate 17.9%, and the remaining four sectors together account for the residual 14.4%. The result is counter-intuitive. Although the trade-credit literature typically emphasizes financial pressure passing through heavy-manufacturing supply chains, in our calibration the largest absolute amplification runs through a sector that figures little in those narratives—simply because it is large enough that an 80% TC share on a -0.8% quarterly shock outweighs a 90% share on a -2.5% shock in a smaller one.

Construction stands between the two patterns. Construction's 76.4% TC share, the lowest among non-RE sectors, reflects its position as the most direct production-network partner of real estate. Its $\text{cum}_B^i = -25.2$ sector %·q— roughly ten times larger in absolute value than the corresponding figure for any other upstream sector—captures the Leontief amplification of real estate's fall in input demand, before any trade credit spills over. The trade-credit channel adds a further -81.3 sector %·q on top of that direct effect, suggesting that the network and credit channels stack roughly additively for the most-exposed upstream sector.

Implications. Three implications follow. First, the modal upstream sector experiences almost none of the regulatory shock through the production network alone—the evidence in Section 3 that exposure to real estate, weighted by each firm's reliance on trade credit, predicts firm-level outcomes is therefore best understood as identifying financial propagation rather than physical input-output propagation. Second, the welfare incidence of the policy is more concentrated than its sectoral footprint: although all seven sectors see TC-mediated declines, three (real estate, services, basic manufacturing) account for 93% of the aggregate GDP loss. Third, the model predicts

a sharp gradient in TC share that future empirical work can test, e.g. by comparing trade credit responses across industries at different Leontief distances from real estate.

The decomposition reported here pools two channels—direct demand through real estate's own contraction and Leontief amplification through input-output linkages—inside the "non-TC" bucket. Separating them would require an additional counterfactual that deactivates the production network itself; we do not pursue this further decomposition because the central finding—that the trade-credit channel accounts for 76–88% of upstream sectoral losses and that the production network on its own transmits very little of the shock past real estate's boundary—is robust to how the residual is split.

5.3 Policy Counterfactuals

The baseline impulse response is conditional on a particular configuration of the regulatory intervention: a sudden one-time tightening of θ_t^H to a peak of -25.9%, decaying through an AR(1) with quarterly persistence of 0.8, and no complementary regulation of firms' trade credit decisions. The actual policy admits several margins of design choice. We use the calibrated model to evaluate three of them: the speed of implementation, the intensity of the underlying tightening, and the presence of a complementary cap on real estate's prepayment behavior. The exercises are intended as positive comparative statics rather than as prescriptive policy advice. Table 12 summarizes the aggregate responses across all three exercises.

5.3.1 Timing: Gradual versus Sudden Implementation

We compare the baseline (sudden) shock to three alternatives in which θ_t^H ramps down linearly over $K \in \{2, 4, 8\}$ quarters before decaying through the same AR(1). The peak deviation is held constant across scenarios so that the long-run policy "stringency" is the same; only the path of implementation differs. Figure 5 reports the GDP impulse responses.

The cumulative 16-quarter GDP loss is approximately invariant to the ramp horizon, but the trough is not monotone. The sudden case has a trough of -4.7% at $t = 1$. The four-quarter ramp produces the shallowest trough, -3.5%. The eight-quarter ramp, however, produces a trough of -5.0% in the middle of the IRF window, slightly deeper than the sudden case. The mechanism

Table 12: Policy counterfactuals: aggregate GDP responses

Scenario	Trough GDP (%, 16q)	Cumulative loss (pp·q, 16q)
Baseline (sudden shock, no complementary regulation)		
$\sigma_{\text{shock}} = 0.30$, $\rho = 0.8$, no cap	-4.70	-30.88
Panel A. Timing: ramp horizon K (peak deviation held fixed)		
$K = 1$ (sudden, baseline)	-4.70	-30.88
$K = 2$	-4.42	-31.05
$K = 4$	-3.52	-31.16
$K = 8$	-5.02	-31.39
Panel B. Intensity: peak shock scaled by s		
$s = 0.5$	-2.52	-15.21
$s = 1.0$ (baseline)	-4.70	-30.88
$s = 1.5$	-6.51	-42.05
Implied non-linearity ($s = 1.5$ vs. linear -46.32)		$+4.27$
Panel C. Trade-credit regulation: floor $\phi^H \geq \phi_{\min}$		
$\phi_{\min}/\bar{\phi}^H = 1.00$ (= Scenario B)	-2.85	-18.24
$\phi_{\min}/\bar{\phi}^H = 0.90$	-2.71	-15.06
$\phi_{\min}/\bar{\phi}^H = 0.75$	-2.55	-11.34
$\phi_{\min}/\bar{\phi}^H = 0.50$	-2.36	-7.40
$\phi_{\min}/\bar{\phi}^H = 0.25$	-2.19	-3.59
$\phi_{\min}/\bar{\phi}^H = 0.00$	-2.05	$+0.04$

Notes. Trough is the deepest quarterly GDP deviation from steady state within the 16-quarter window. Cumulative loss is the sum of quarterly GDP deviations over the same window. The baseline row reproduces the model's response under the calibrated shock and no complementary regulation; it is repeated as the reference point for each panel. Panel A varies the number of quarters K over which θ^H is linearly phased down to its trough value, holding the peak deviation fixed. Panel B scales the peak deviation by s, holding the shock shape fixed. Panel C imposes a lower bound on real estate's prepayment fraction ϕ^H, parameterized as a fraction of the steady-state value $\bar{\phi}^H$; the $\phi_{\min}/\bar{\phi}^H = 1.00$ row reproduces the fixed-ϕ counterfactual of Section 5.1 (Scenario B). Numbers reported to two decimals; values may shift slightly across runs because of fsolve convergence tolerance.

is the χ_t-wedge: with $\rho_\chi = 0.92$, the demand-side wedge is slow-moving and continues to widen even after θ_t^H begins to revert. Under a sufficiently slow rollout the two slow variables (θ_t^H approaching its trough; χ_t continuing to fall) align at mid-horizon, deepening the worst quarter. The result suggests that spreading a credit-policy shock over many quarters does not reduce the maximum quarterly impact in our calibration and can amplify it through the slower demand-side channel.

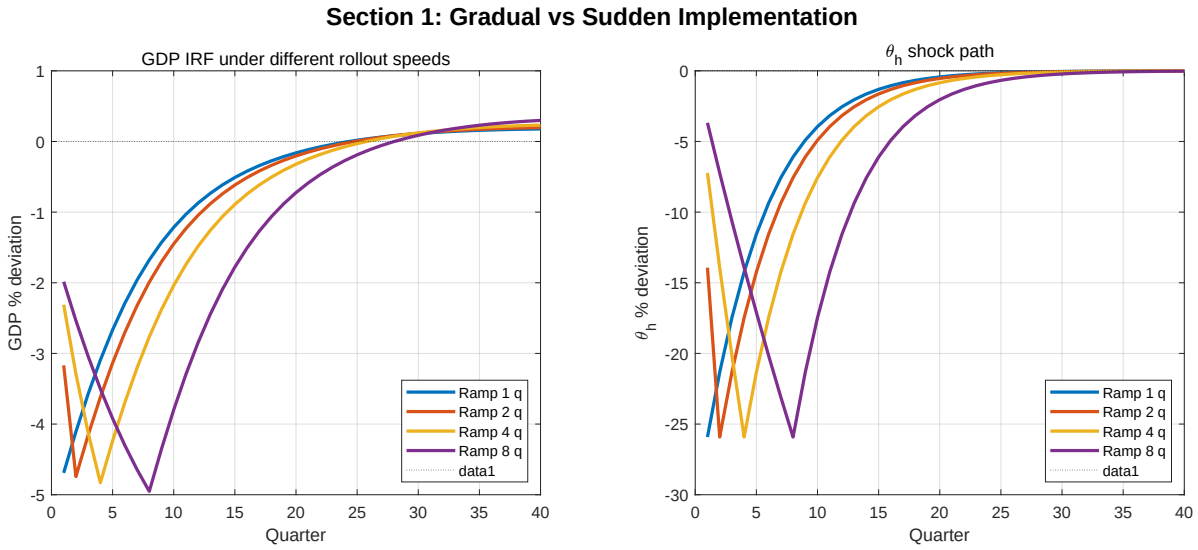


Figure 5

Notes: GDP IRF and shock path under four implementation speeds. Ramp K refers to the number of quarters over which θ_t^H is linearly phased down to its trough value before AR(1) recovery sets in. The peak deviation is held constant at -25.9%.

5.3.2 Intensity Scaling

We rescale the peak magnitude of the shock by factors 0.5 and 1.5 relative to the baseline. Figure 6 reports the cumulative GDP loss as a function of the scaling factor. The response is close to linear at the baseline scale but exhibits modest sub-linearity at higher intensities: a $1.5\times$ shock produces a cumulative loss of approximately -42 pp·q, against a linear extrapolation from the baseline of -45 pp·q. The departure from linearity reflects a partial saturation of the borrowing-constraint multiplier μ_t^H: once the trade-credit channel is fully engaged, additional tightening produces diminishing amplification at the margin. Over the range we consider, the trade-credit channel is approximately—though not exactly—linear in the magnitude of the shock.

5.3.3 Trade-Credit Regulation

The third exercise asks whether a complementary regulation of real estate's trade-credit behavior can mitigate the upstream spillover identified in Section 5.2. We introduce a floor $\phi_{\min}$ on real estate's prepayment fraction, requiring $\phi_t^H \geq \phi_{\min}$ in every period. The unregulated baseline corresponds to a non-binding floor. We sweep six levels of $\phi_{\min}$, parameterized as fractions $\rho \in$

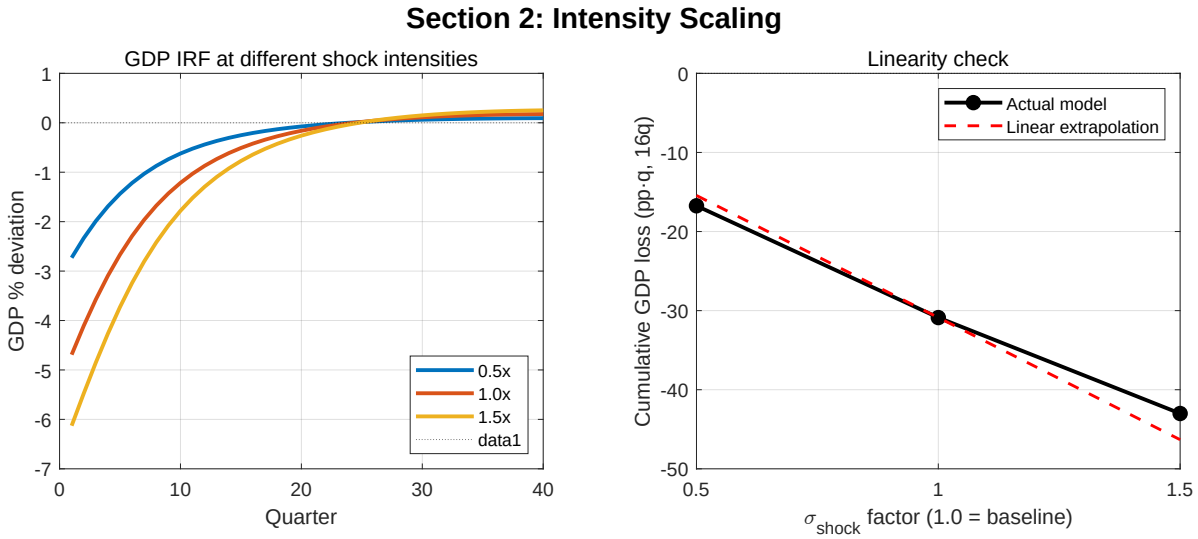


Figure 6

Notes: GDP IRF at three intensities of the underlying credit shock, and cumulative GDP loss as a function of shock magnitude. The right panel compares the model-implied cum loss (black) to a linear extrapolation from the baseline (dashed red).

$\{0, 0.25, 0.5, 0.75, 0.9, 1\}$ of the steady-state value $\bar{\phi}^H$.

Before turning to the results, we note the steady-state level of ϕ^H. In the calibrated steady state, the real estate sector's borrowing-constraint multiplier $\bar{\mu}^H$ is large (approximately 1.03), which through the first-order condition $\phi^H = 1 - R^H - \mu^H$ implies $\bar{\phi}^H \approx -0.08$. A negative $\bar{\phi}^H$ does not have a literal interpretation as a prepayment fraction in $[0, 1]$; rather, it indicates that, in the private equilibrium, real estate extracts net trade credit from its suppliers on a chronic basis—paying less cash upfront than the zero-trade-credit benchmark would imply. This feature of the calibration is consistent with the empirical pattern that large Chinese real estate developers maintain elevated accounts-payable ratios and routinely defer payments to upstream suppliers even outside of credit-tightening episodes, as documented in the descriptive statistics for the real estate sector in Section 3 and in widely reported industry practice. The policy exercise in this subsection should be read against this background: the unregulated steady state already exhibits substantial extraction of trade credit by real estate, and the shock amplifies it.

The cap at $\rho = 1$ corresponds to $\phi_{\min} = \bar{\phi}^H$. This case prevents the shock-induced reduction of ϕ_t^H but does not alter its steady-state level; it is equivalent to the fixed-ϕ counterfactual of Section 5.1. The cap at $\rho = 0$ corresponds to $\phi_{\min} = 0$, which represents a stronger structural

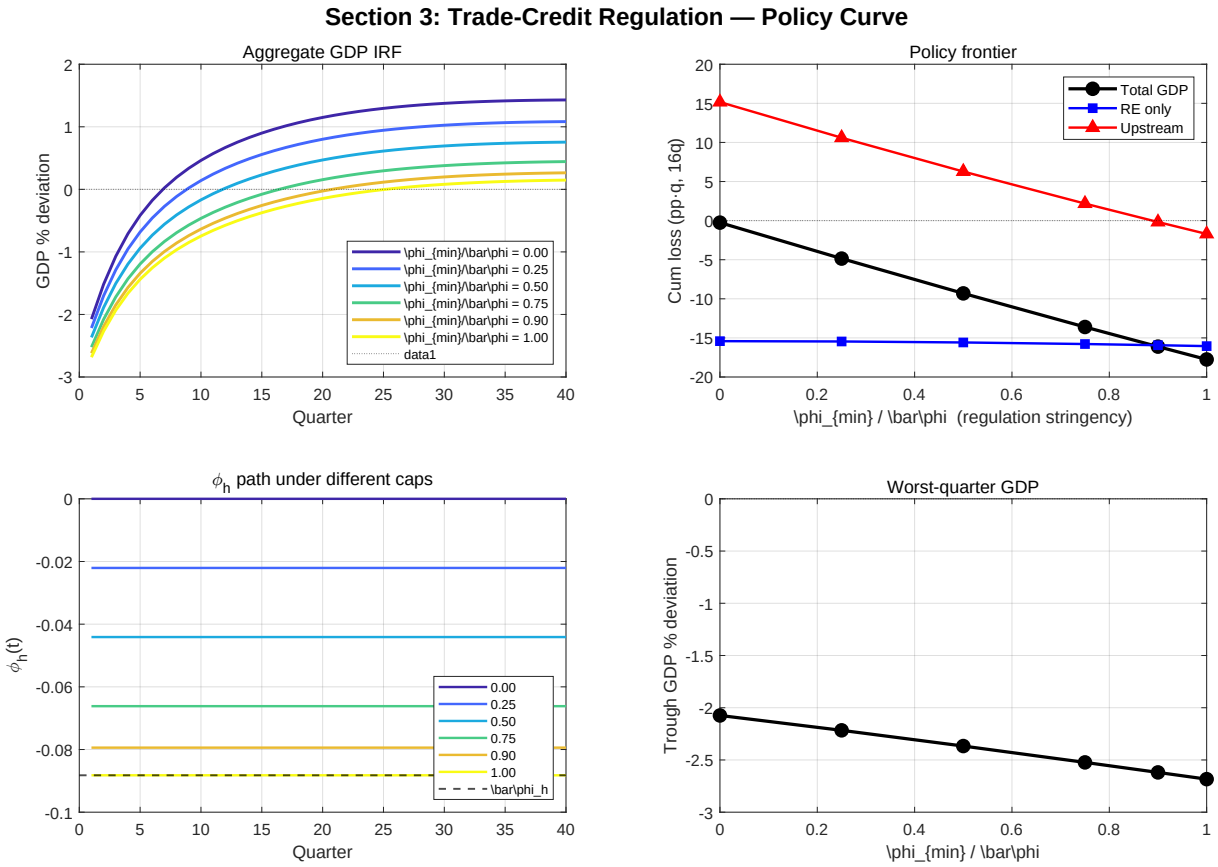


Figure 7

Notes: Policy curve under trade-credit regulation. *Top-left:* aggregate GDP IRF under each $\phi_{\min}$. *Top-right:* cumulative 16-quarter loss decomposed into real estate's contribution and the aggregate non-RE contribution. *Bottom-left:* the path of ϕ_t^H under each cap. *Bottom-right:* worst-quarter GDP deviation.

intervention: it requires real estate to pay suppliers at least at par, ruling out the steady-state extraction of trade credit. For $\rho < 1$, the cap is binding in steady state, and the exercise should be read as the joint response to (i) the Three Red Lines shock and (ii) a contemporaneous, permanent regulation of payment terms.

Three patterns emerge from Figure 7.

The aggregate GDP loss is monotone and approximately linear in the cap stringency. As ρ falls from 1 to 0, the cumulative 16-quarter GDP loss declines from approximately -18 pp·q to approximately 0, traversing a roughly linear interior. The value at $\rho = 1$ (-18.2 pp·q) reproduces the fixed-ϕ counterfactual of Section 5.1, providing an internal consistency check. The worst-quarter GDP deviation is similarly attenuated by the cap, from approximately -2.8% at $\rho = 1$ to

-2.0% at $\rho = 0$, against the -4.7% trough of the unregulated baseline.

Real estate's contribution to the aggregate is relatively invariant to the cap. Across all six values of ρ, real estate's cumulative contribution to GDP lies in the range -15 to -17 pp·q. The cap forces real estate to bear more of the financing pressure on its own balance sheet, but in general equilibrium the upstream sector's improved liquidity reduces input costs and partially offsets the direct burden. The net effect on real estate is small relative to the variation in upstream outcomes.

The upstream sectors absorb the entire variation along the policy curve. The upstream cumulative contribution ranges from approximately -2 pp·q at $\rho = 1$ to approximately $+15$ pp·q at $\rho = 0$. The mechanism is straightforward: the cap forces real estate to retain financing pressure rather than transmit it through the trade-credit chain, so upstream sectors retain—and at the strictest cap, expand—their liquidity buffer relative to the unregulated case.

The exercise should not be over-interpreted as a policy prescription. The cap does not provide real estate with additional resources; it forecloses the substitution into trade credit and thereby forces real estate to absorb a larger share of the financing pressure on its own balance sheet through reduced output. The model treats this adjustment as smooth, but in practice the strictest caps in our sweep would likely push real estate into a corner of discrete default and project incompletion—phenomena that were material features of the 2021–2024 episode and that lie outside the model's coverage. Once such discrete events are taken into account, the losses generated by very strict caps would partially reverse through write-downs of accounts receivable on the upstream side. The policy curve in Figure 7 therefore identifies a *direction* (binding the trade-credit channel reduces aggregate spillover) and a rough *upper bound* on the GDP-loss reduction, rather than a calibrated optimum.

A broader lesson runs through the three exercises. The aggregate loss from the Three Red Lines policy is largely insensitive to the timing and intensity of the underlying shock once peak stringency is held fixed: spreading the shock over more quarters does not reduce the maximum quarterly impact, and scaling the shock down lowers the spillover roughly in proportion to the reduction in regulatory bite. What materially shapes the aggregate response is the configuration of the trade-credit chain through which the shock propagates. The common thread is that, for

a sector as deeply embedded in the production network as real estate, regulatory interventions cannot be evaluated solely on the basis of their direct effect on the targeted firms; the spillover through endogenous trade credit can exceed the direct loss in magnitude and should be built into the policy's design.

5.4 Robustness of the Quantitative Findings

Three considerations support reading the headline numbers as properties of the mechanism rather than artifacts of particular modeling choices.

First, the decomposition is invariant to the comparison window. Computed over the 20-quarter window that matches the regression sample, the trade-credit share of the cumulative GDP loss is approximately 40% (-31.9 versus -19.1 pp·q); computed over the 16-quarter window used in the sectoral decomposition of Section 5.2, it is 40.9%. The EILR comparison, in contrast, is window-sensitive—the model's response is front-loaded while the empirical effect is persistent—which is exactly why we headline the decomposition and read the EILR only as a check on the order of magnitude.

Second, the share does not lean on the AR(1) shape of the shock. The decomposition holds the path of θ_t^H—and the induced path of the demand wedge χ_t—fixed across Scenarios A and B; the only difference between the scenarios is whether the prepayment fraction responds to the borrowing constraint. The roughly 40% share is therefore a statement about the mechanism conditional on the shock, not about the assumed shock process. The shape of the path does matter for the level response—the timing exercises of Section 5.3 show that ramped implementations move the depth and timing of the GDP trough—but the attribution to trade credit is computed with the path held fixed.

Third, the housing demand wedge does not contaminate the trade-credit numbers. As discussed in Section 4.2.4, χ_t enters only the household housing Euler equation; the steady-state trade-credit block—the multipliers μ_i and the prepayment fractions ϕ_i—is invariant to the level of $\bar{\chi}$. And because the decomposition holds the path of χ_t fixed across Scenarios A and B, the wedge contributes identically to both scenarios and cancels from the comparison: the roughly 40% share is unaffected by it. In sum, the decomposition survives the three probes a skeptic would apply

first—the window, the shock path, and the auxiliary demand wedge.

6 Conclusion

This paper asks how a leverage cap imposed on a single sector travels through the rest of the economy, and answers that it travels along the network's own transaction edges, carried by the credit written on them. Exploiting China's 2020 "Three Red Lines" policy as a quasi-natural experiment, we document that firms more exposed to real estate through trade credit experienced significant declines in cash holdings, sales, borrowing, and equity values, with effects emerging sharply after the policy and persisting for over three years. Evidence from both sides of the trade-credit relationship—accounts payable rising at constrained developers and accounts receivable rising at their suppliers—identifies trade credit as the operative channel of transmission: what appears as delayed payment on one balance sheet appears as involuntary lending on another. To interpret and quantify this mechanism, we develop a dynamic production-network model in which each buyer's prepayment fraction is pinned down by the shadow value of its borrowing constraint, so that payment terms respond endogenously to financial conditions. Calibrated to China's 2020 input–output table, the model generates a GDP trough of -4.7%, and a counterfactual decomposition attributes roughly 40% of the cumulative GDP impact to the endogenous adjustment of trade credit—a $1.7\times$ amplification relative to a model that holds payment terms fixed. An independent aggregation of the reduced-form estimates implies an economy-wide loss of the same order of magnitude as the model's prediction; we read this convergence as an over-identifying check on the mechanism rather than as a precise match.

The central insight is that trade credit transforms what would otherwise be a localized regulatory intervention into an economy-wide disturbance. When credit-constrained firms delay payments to their suppliers, they convert their own liquidity shortage into a contraction of credit supply for their trading partners. In dense production networks, this mechanism propagates distress far beyond the targeted sector—through linkages that standard input-output analysis treats as purely physical but that, in the data and in our model, carry a financial friction that is quantitatively first-order. Our framework makes this transmission tractable: the same logic that operates between any pair of constrained firms aggregates into the network amplification we document.

These findings carry implications for both policymakers and researchers. For policymakers, the results suggest that when a sector is as deeply embedded in production networks as Chinese real estate, the transmission costs of sector-specific financial regulation operate primarily through the trade-credit chain rather than through the direct effect on the targeted firms. Spillovers of this kind should be built into the design of any sector-specific intervention, and assessments that focus only on the targeted firms will substantially underestimate the aggregate response. We measure only the cost side of the ledger: the policy may well have purchased durable benefits—lower developer leverage, reduced systemic risk—that our framework does not model. The lesson is therefore not a verdict on the Three Red Lines but a requirement for reaching one: the cost carried along payment terms belongs in the accounting. For researchers, our framework provides a portable tool for studying other settings where sector-specific financial shocks meet production networks—banking regulation, sector-targeted industrial policy, and credit market interventions in other emerging economies. The methodology of combining firm-level reliance on trade credit with network-derived exposure measures, developed here to overcome the scarcity of data on bilateral credit flows, applies wherever such bilateral data are unavailable.

We close with two caveats. First, the multi-sector application adapts the framework to the housing setting in two respects: the durable-asset structure of housing, and the time-varying preference wedge χ_t that captures households' demand-side response. The latter is reduced-form, calibrated to match the observed housing-price decline during 2021–2024; a more complete treatment would micro-found it through explicit mortgage markets and household heterogeneity in housing tenure. Both elements are housing-specific adaptations rather than features of the underlying trade-credit framework: the core mechanism—endogenous trade credit responding to binding borrowing constraints in a dynamic production network—does not depend on them. Second, our empirical analysis is restricted to listed firms, which may underrepresent the small and medium enterprises most likely to be affected by shocks to trade credit; richer firm-level data could reveal heterogeneous effects across the distribution of firm size.

Several extensions of this work are worth pursuing. The framework can be applied to other episodes of sector-specific financial regulation—deleveraging in banking, restrictions on the energy sector, or trade policy shocks—to assess whether amplification through trade credit is similarly important elsewhere. We could enrich the model's endogenous trade credit by adding

heterogeneous default risk, multi-period trade-credit contracts, or explicit factoring and supply-chain-finance instruments. More broadly, our results suggest that interfirm credit networks deserve a more central place in macroeconomic analysis: they are pervasive, quantitatively significant, and respond endogenously to policy in ways that standard production-network or financial-accelerator frameworks do not capture. Building richer models of these networks—and gathering the bilateral firm-pair data needed to discipline them—is a promising direction for future research. The lesson for macroprudential design, however, stands already. Han et al. (2021) show that a targeted financial constraint reshapes strategic behavior within the market it regulates; our results carry the point across markets: a leverage cap aimed at one sector is, through the credit written on the network's own transaction edges, a policy for every sector that trades with it—and should be evaluated as such.

References

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "Systemic Risk and Stability in Financial Networks," *American Economic Review*, 2015, *105* (2), 564–608.

_ , Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, "The Network Origins of Aggregate Fluctuations," *Econometrica*, 2012, *80* (5), 1977–2016.

Acharya, Viral V., Katharina Bergant, Matteo Crosignani, Tim Eisert, and Fergal McCann, "The Anatomy of the Transmission of Macroprudential Policies," *Journal of Finance*, 2022, *77* (5), 2533–2575.

Adelino, Manuel, Miguel A. Ferreira, Mariassunta Giannetti, and Pedro Pires, "Trade Credit and the Transmission of Unconventional Monetary Policy," *Review of Financial Studies*, 2023, *36* (2), 775–813.

Almeida, Heitor, Daniel Carvalho, and Taehyun Kim, "The Working Capital Credit Multiplier," *Journal of Finance*, 2024, *79* (6), 4247–4302.

Altinoglu, Levent, "The Origins of Aggregate Fluctuations in a Credit Network Economy," *Journal of Monetary Economics*, 2021, *117*, 316–334.

Baqaee, David Rezza and Emmanuel Farhi, "The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem," *Econometrica*, 2019, *87* (4), 1155–1203.

Boissay, Frédéric and Reint Gropp, "Payment Defaults and Interfirm Liquidity Provision," *Review of Finance*, 2013, *17* (6), 1853–1894.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna, "Difference-in-Differences with a Continuous Treatment," NBER Working Paper 32117, National Bureau of Economic Research 2024.

Carvalho, Vasco M., "From Micro to Macro via Production Networks," *Journal of Economic Perspectives*, 2014, *28* (4), 23–48.

Chen, Kaiji, Huancheng Du, and Chang Ma, "The Spillover Effects of Real Estate," 2024. Working paper, presented at ABFER 2024 Annual Conference and AEA 2025 Meetings.

Costello, Anna M., "Credit Market Disruptions and Liquidity Spillover Effects in the Supply Chain," *Journal of Political Economy*, 2020, *128* (9), 3434–3468.

Cun, Wukuang, Vincenzo Quadrini, Qi Sun, and Junjie Xia, "Dynamics of Trade Credit in China," *Economic Journal*, 2022, *132* (648), 2702–2736.

Cuñat, Vicente, "Trade Credit: Suppliers as Debt Collectors and Insurance Providers," *Review of Financial Studies*, 2007, *20* (2), 491–527.

DeFusco, Anthony A., Stephanie Johnson, and John Mondragon, "Regulating Household Leverage," *Review of Economic Studies*, 2020, *87* (2), 914–958.

Donaldson, Paula, "Cross-sectional Identification with Heterogeneous Exposure to General Equilibrium Effects," Job Market Paper, University of California, San Diego 2026.

Garcia-Appendini, Emilia and Judit Montoriol-Garriga, "Firms as Liquidity Providers: Evidence from the 2007–2008 Financial Crisis," *Journal of Financial Economics*, 2013, *109* (1), 272–291.

Han, Lu, Chandler Lutz, Benjamin Sand, and Derek Stacey, "The Effects of a Targeted Financial Constraint on the Housing Market," *Review of Financial Studies*, 2021, *34* (8), 3742–3788.

He, Jianyao, Dun Jia, Kai Li, and Wenbin Wu, "A High-Frequency Measure of Chinese Monetary Policy Shocks," Working Paper, Peking University HSBC Business School 2026.

Huremović, Kenan, Gabriel Jiménez, Enrique Moral-Benito, José-Luis Peydró, and Fernando Vega-Redondo, "Production and Financial Networks in Interplay: Crisis Evidence from Supplier-Customer and Credit Registers," *American Economic Review*, 2026.

Iacoviello, Matteo, "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle," *American Economic Review*, 2005, *95* (3), 739–764.

_ **and Stefano Neri**, "Housing Market Spillovers: Evidence from an Estimated DSGE Model," *American Economic Journal: Macroeconomics*, 2010, *2* (2), 125–164.

Jacobson, Tor and Erik von Schedvin, "Trade Credit and the Propagation of Corporate Failure: An Empirical Analysis," *Econometrica*, 2015, *83* (4), 1315–1371.

Jiménez, Gabriel, Steven Ongena, José-Luis Peydró, and Jesús Saurina, "Macroprudential Policy, Countercyclical Bank Capital Buffers, and Credit Supply: Evidence from the Spanish Dynamic Provisioning Experiments," *Journal of Political Economy*, 2017, *125* (6), 2126–2177.

Liu, Zheng, Pengfei Wang, and Tao Zha, "Land-Price Dynamics and Macroeconomic Fluctuations," *Econometrica*, 2013, *81* (3), 1147–1184.

Long, John B. and Charles I. Plosser, "Real Business Cycles," *Journal of Political Economy*, 1983, *91* (1), 39–69.

Luo, Shaowen, "Propagation of Financial Shocks in an Input-Output Economy with Trade and Financial Linkages of Firms," *Review of Economic Dynamics*, 2020, *36*, 246–269.

Ma, Chang and Yong Wang, "Real Estate Policy Regulation and Corporate Financial Risk: China's Three Red Lines," *Pacific Economic Review*, 2025. Forthcoming.

Murfin, Justin and Ken Njoroge, "The Implicit Costs of Trade Credit Borrowing by Large Firms," *Review of Financial Studies*, 2015, *28* (1), 112–145.

Ozdagli, Ali and Michael Weber, "Monetary Policy through Production Networks: Evidence from the Stock Market," *Review of Financial Studies*, 2026, *39* (5), 1411–1462.

People's Bank of China, "2019 Urban Chinese Household Asset and Liability Survey," Survey report, Financial Stability Bureau 2020.

Petersen, Mitchell A. and Raghuram G. Rajan, "Trade Credit: Theories and Evidence," *Review of Financial Studies*, 1997, *10* (3), 661–691.

Reischer, Marvin, "Finance-Thy-Neighbor: Trade Credit Origins of Aggregate Fluctuations," 2020. Working paper.

Rogoff, Kenneth S. and Yuanchen Yang, "Has China's Housing Production Peaked?," *China & World Economy*, 2021, *29* (1), 1–31.

A Empirical Robustness and Extensions

This appendix collects auxiliary empirical analyses underlying the findings of Section 3. The continuous dose-response specification (A.3) verifies that the channel operates monotonically across the intensity distribution rather than only at extreme exposure. The direct trade-credit dynamics (A.4) document the payment-delay mechanism in the data. The real-economy outcomes (A.5) trace the channel beyond balance sheets to investment, inventory, and profitability. The financial-constraint heterogeneity (A.6) provides the full set of outcomes and constraint proxies that motivate the joint test in Section 3.6. The remaining sections report the full network decomposition across outcomes (A.7), compare the ecological treatment intensity to a direct dyadic measure constructed from Wind disclosure data (A.8), verify that the economy-wide estimate is not an artifact of differential sensitivity to aggregate financial conditions (A.1), and show that it is robust to dropping any single industry and to two-way clustering (A.2).

A.1 Sensitivity to Aggregate Financial Conditions

Our economy-wide design regresses firm outcomes on a cross-sectional exposure measure while absorbing aggregate movements with industry-by-time fixed effects. Because those effects absorb the common component of *every* aggregate shock—named or not—the concern is not an open-ended list of controls to be added one by one, but a single narrow channel the fixed effects cannot reach: as Donaldson (2026) shows, a design of this kind identifies the partial-equilibrium object only if exposure to the shock of interest is uncorrelated with the *differential* incidence, across firms, of other aggregates that move in general equilibrium. Her diagnostic, which we adopt, tests for exactly this—it adds the interaction of the exposure measure with a suspect aggregate and asks whether the coefficient of interest is stable. A stable coefficient means the estimate does not depend on that channel being absent, rather than needing the interaction to rescue it. The dominant suspect in our setting is monetary policy, which eased sharply as the deleveraging bit, and we stress-test the estimate against it below using both the observed conditions and a properly identified exogenous shock. Where a design feature already closes a channel—the within-developer comparison holds housing demand fixed, and the incidence runs through higher-order network paths rather than direct sales to real estate—we rely on that rather than on further interactions.

The sharpest form of the concern is that higher-exposure firms load differentially on *exogenous* monetary policy surprises. We use the high-frequency China monetary policy shock of He et al. (2026), identified from the response of negotiable-certificate-of-deposit yields around policy announcements and aggregated to the quarterly surprise, and add its interaction with Intensity to equation (2), with the shock's own level absorbed by the time effects. Table 13 shows the coefficient of interest is essentially unchanged: on the same sample it moves from -2.41 to -2.47 for cash and from -2.13 to -2.11 for sales without the size control, and is equally stable with it. Higher-exposure firms show a modest but statistically significant differential loading on monetary surprises in their cash holdings, but none in sales, and in neither case does controlling for it materially move the estimate. Because this shock is exogenous by construction, it meets the concern more directly than the observed conditions we turn to next.

Table 13: Robustness to an Exogenous High-Frequency Monetary Policy Shock

	Cash		Sales	
	(1)	(2)	(3)	(4)
Intensity(%) × Post	-2.471***	-3.630***	-2.108**	-3.507***
	(0.488)	(0.553)	(0.831)	(1.028)
Intensity × MP shock	0.205***	0.209***	-0.061	-0.076
	(0.038)	(0.039)	(0.048)	(0.050)
N	127430	104804	124947	103856
Firm FE	Yes	Yes	Yes	Yes
Industry × Time FE	Yes	Yes	Yes	Yes
Pre-policy size × Post	No	Yes	No	Yes
Adj. R^2	0.904	0.905	0.949	0.952

Notes. Each column re-estimates equation (2) adding the interaction of Intensity with the high-frequency China monetary policy shock of He et al. (2026) (the quarterly surprise; its own level is absorbed by the time effects). Columns (1)–(2) are cash holdings, (3)–(4) sales; columns (2) and (4) add the pre-policy size control. On the same sample, the coefficient of interest *without* the interaction is -2.41, -3.57, -2.13, and -3.53—so adding the exogenous monetary shock leaves it essentially unchanged. Standard errors clustered at the industry level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

We complement this with a broader battery that uses the observed levels of eight aggregate conditions. We re-estimate equation (2) adding, one at a time, the interaction of Intensity with a measure of aggregate financial conditions; the measure's own level is absorbed by the quarter effects, so the interaction identifies whether higher-exposure firms load differentially on that aggregate. Table 14 reports six representative measures spanning prices (the seven-day interbank rate DR007 and the five-year loan prime rate) and quantities (the growth of total social financing and of

M2, and the reserve requirement ratio); Table 15 adds three further money-market rates—the all-institution repo rate R007, three-month SHIBOR, and the seven-day open-market-operation policy rate, and reports the coefficient of interest under both the parsimonious and the size-controlled specifications (Panels A and B). It stays negative and statistically significant under every one of the eight measures, for both outcomes and both specifications; the size-controlled estimates are, as elsewhere, uniformly larger. The sales effect is essentially unchanged throughout; the cash effect attenuates modestly under the rate and reserve measures but never falls below roughly 160 million RMB per percentage point of exposure. Because these aggregates co-move strongly over 2020–2022—the money-market rates are pairwise correlated above 0.9—we add them one at a time rather than jointly: entered together, their interactions with Intensity are collinear and individually uninformative. Even so, including all eight interactions at once leaves the coefficient of interest negative and significant (-1.46 for cash, -1.45 for sales), as does a parsimonious rate–quantity–reserve combination (-1.33 and -1.77); the effect attenuates but does not disappear when every differential loading is absorbed simultaneously.

Panel B of Table 2 showed that exposure partly proxies for firm size, and so does the modest cash attenuation just described: the Panel B columns of Table 15 add the predetermined size control, under which the cash coefficient is uniformly larger and the attenuation shrinks—from twenty-eight to twenty-two percent under DR007, and from thirty-three to twenty-seven percent under the reserve requirement—so part of the apparent monetary sensitivity of cash was a size effect rather than a channel confound.

A.2 Influence and Two-Way Inference

Two further checks confirm that the economy-wide estimate is neither driven by a single sector nor an artifact of the inference. Figure 8 re-estimates the cash and sales regressions dropping one industry at a time. In the baseline (top row), for both outcomes every one of the seventy-six leave-one-out coefficients is negative, and all but one lie in a narrow band around the full-sample estimate (dashed line): the cash coefficient is significant in all seventy-six fits and the sales coefficient in seventy-five. The single exception, and the only estimate that moves materially, is the fit that drops real estate itself. Removing the directly shocked sector makes the coefficient *larger*

Table 14: Robustness to Differential Sensitivity to Aggregate Financial Conditions

	(1) Baseline	(2) +DR007	(3) +LPR 5Y	(4) +TSF	(5) +M2	(6) +RRR
Panel A: Cash holdings (100M RMB)						
Intensity(%) × Post	-2.362***	-1.712***	-1.587***	-2.186***	-2.306***	-1.574***
	(0.536)	(0.348)	(0.316)	(0.404)	(0.535)	(0.329)
× DR007		1.108***				
		(0.338)				
× LPR 5Y			-0.801			
			(0.584)			
× TSF growth				0.067		
				(0.059)		
× M2 growth					-0.177***	
					(0.021)	
× RRR						0.212***
						(0.066)
N	142024	142024	109437	142024	142024	142024
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Pre-policy size × Post	No	No	No	No	No	No
Adj. R^2	0.899	0.899	0.924	0.899	0.899	0.899
Panel B: Sales (100M RMB)						
Intensity(%) × Post	-2.037**	-1.814***	-1.814***	-1.907***	-2.006**	-1.898***
	(0.800)	(0.585)	(0.642)	(0.716)	(0.785)	(0.639)
× DR007		0.381				
		(0.401)				
× LPR 5Y			-0.284			
			(0.266)			
× TSF growth				0.048		
				(0.049)		
× M2 growth					-0.105*	
					(0.058)	
× RRR						0.037
						(0.073)
N	139429	139429	107398	139429	139429	139429
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Pre-policy size × Post	No	No	No	No	No	No
Adj. R^2	0.951	0.951	0.965	0.951	0.951	0.951

Notes. Each column adds to the baseline the interaction of Intensity with one measure of aggregate financial conditions (its own level absorbed by the time effects); the indented rows report that interaction's coefficient. DR007: seven-day interbank repo rate; LPR 5Y: five-year loan prime rate; TSF and M2: year-on-year growth of total social financing and M2; RRR: reserve requirement ratio. Standard errors clustered at the industry level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 15: Robustness to Aggregate Financial Conditions: All Eight Measures

	Cash		Sales	
	Panel A	Panel B	Panel A	Panel B
Baseline	-2.362***	-3.530***	-2.037**	-3.418***
	(0.536)	(0.611)	(0.800)	(0.997)
DR007	-1.712***	-2.762***	-1.814***	-3.261***
	(0.348)	(0.458)	(0.585)	(0.804)
R007	-1.755***	-2.832***	-2.070***	-3.520***
	(0.416)	(0.516)	(0.604)	(0.830)
SHIBOR 3M	-1.688***	-2.776***	-1.963***	-3.416***
	(0.460)	(0.557)	(0.623)	(0.843)
OMO 7d	-2.582***	-3.690***	-2.163**	-3.562***
	(0.323)	(0.430)	(0.824)	(1.007)
LPR 5Y	-1.587***	-2.578***	-1.814***	-3.047***
	(0.316)	(0.397)	(0.642)	(0.817)
TSF growth	-2.186***	-3.289***	-1.907***	-3.345***
	(0.404)	(0.497)	(0.716)	(0.920)
M2 growth	-2.306***	-3.468***	-2.006**	-3.388***
	(0.535)	(0.612)	(0.785)	(0.985)
RRR	-1.574***	-2.592***	-1.898***	-3.361***
	(0.329)	(0.432)	(0.639)	(0.838)

Notes. Each row adds the interaction of Intensity with one measure of aggregate financial conditions; the entry is the coefficient on Intensity×Post—the coefficient of interest. Panel A columns use the parsimonious specification; Panel B columns add the pre-policy size control (log pre-policy employment × Post). Measures span money-market rates (DR007, R007, three-month SHIBOR, the seven-day OMO policy rate), the five-year LPR, and credit quantities (TSF and M2 growth, the reserve requirement ratio). All specifications include firm, quarter, and industry-by-time fixed effects. N = 142,024/139,429 (Panel A, cash/sales) and 114,771/113,783 (Panel B); the LPR 5Y rows use fewer observations (the series begins in 2019). Standard errors clustered at the industry level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

in magnitude—from -2.36 to -3.60 for cash and from -2.04 to -3.20 for sales—while widening its confidence interval, because real estate firms, as the policy's direct target, dilute the upstream spillover they help cause. No industry's removal weakens the result; the one sector a reader might suspect of driving it works, if anything, against it. The bottom row of Figure 8 repeats the exercise with the pre-policy size control: the pattern is identical and the coefficients are uniformly larger (full-sample estimates of -3.53 for cash and -3.42 for sales), with the real-estate fit again the sole outlier. Table 16 recomputes the economy-wide standard errors clustering simultaneously by industry and by quarter, which allows a common aggregate shock to correlate outcomes across firms within a period. Under both the parsimonious and the size-controlled specifications (Panels A and B), the two-way standard errors differ from the industry-clustered baseline by less than three percent, and all four coefficients remain significant at conventional levels.

Table 16: Economy-Wide Estimates under Two-Way Clustering

	Dependent variable (100M RMB)			
	(1) Cash	(2) Sales	(3) Debt	(4) Equity
Panel A: baseline				
Intensity(%) × Post	-2.362***	-2.037**	-2.246***	-3.623***
	(0.534)	(0.806)	(0.700)	(0.867)
N	142024	139429	142024	142024
Firm FE	Yes	Yes	Yes	Yes
Industry × Time FE	Yes	Yes	Yes	Yes
Pre-policy size × Post	No	No	No	No
Adj. R^2	0.899	0.951	0.897	0.976
Panel B: with pre-policy size control				
Intensity(%) × Post	-3.530***	-3.418***	-3.802***	-5.779***
	(0.615)	(1.027)	(1.097)	(0.961)
ln(Emp, pre) × Post	14.326***	19.644***	18.993**	29.793***
	(3.204)	(6.223)	(9.174)	(6.584)
N	114771	113783	114771	114771
Firm FE	Yes	Yes	Yes	Yes
Industry × Time FE	Yes	Yes	Yes	Yes
Pre-policy size × Post	Yes	Yes	Yes	Yes
Adj. R^2	0.899	0.953	0.898	0.973

Notes. The four economy-wide regressions of Table 2 with standard errors clustered simultaneously by industry and by quarter. Panel A is the parsimonious specification; Panel B adds the pre-policy size control (log pre-policy employment × Post). Within each panel the point estimates equal the industry-clustered baseline; the two-way standard errors, in parentheses, differ from the industry-clustered ones by less than three percent. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

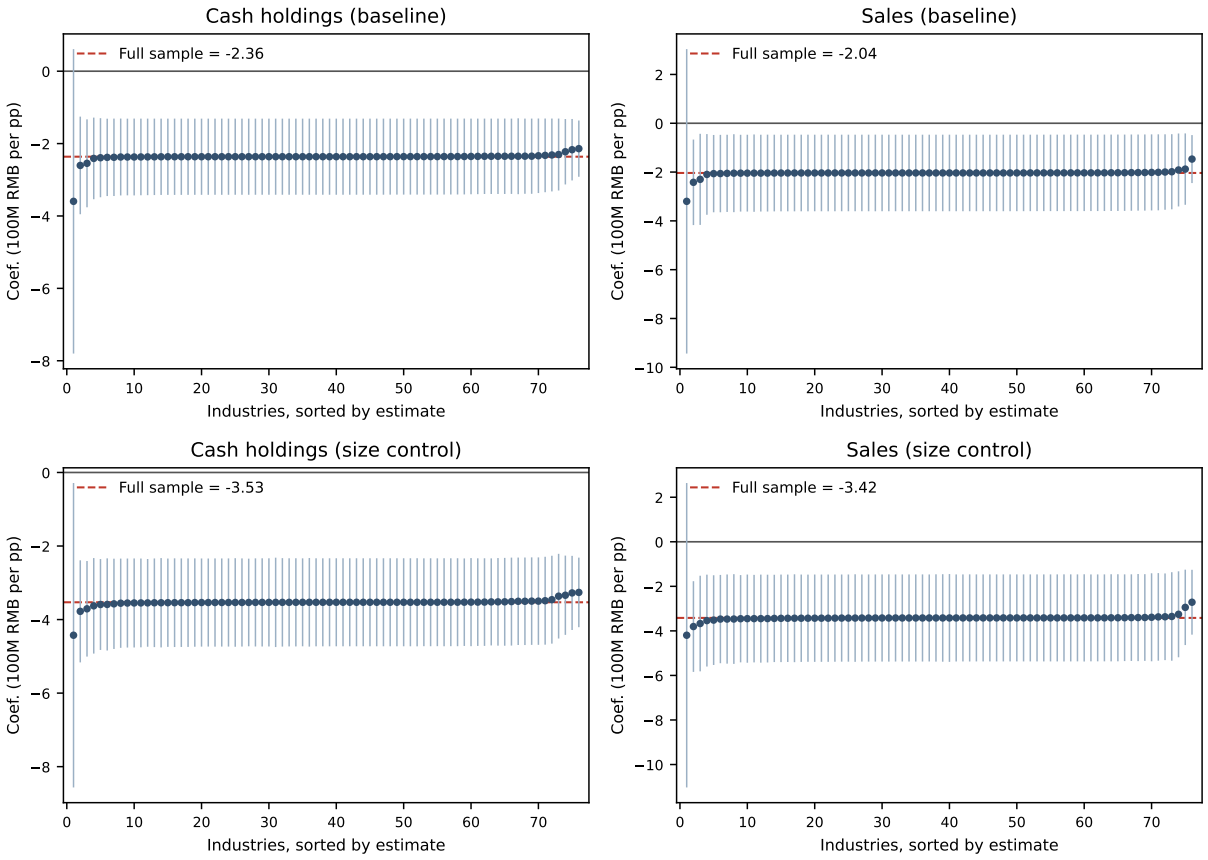


Figure 8: Leave-One-Industry-Out Estimates of the Economy-Wide Spillover

Notes. Each dot re-estimates equation (2) with one industry removed (seventy-six industries in all), sorted along the horizontal axis by the point estimate; whiskers are 95% confidence intervals from industry-clustered standard errors. The dashed line is the full-sample estimate; the solid line is zero. Top row: the parsimonious specification; bottom row: with the pre-policy size control. Left column: cash holdings; right column: sales; both in 100M RMB per percentage point of exposure.

A.3 Dose-Response and Heterogeneous Exposure

Callaway et al. (2024) note that the continuous-treatment difference-in-differences slope identifies a weighted average of dose-specific level effects under standard parallel trends, but recovering a marginal causal response to dose requires the stronger assumption that high-exposure firms would have responded identically to lower exposure—ruling out selection into dose and treatment-effect heterogeneity across exposure types. We therefore do not interpret β in our main specification as a structural marginal effect of an infinitesimal change in exposure. Instead, we present two complementary forms of evidence that the heterogeneous post-policy decline aligns with pre-policy exposure in a dose-response pattern.

Quintile dose-response. Sorting firms into quintiles of pre-policy exposure intensity, we estimate, with the lowest-exposure quintile Q_1 as the reference,

$$Y_{it} \;=\; \sum_{k=2}^{5} \beta_k \cdot \big(\mathbf{1}\{Q_i = k\} \times \text{Post}_t\big) \;+\; \theta_i \;+\; \gamma_{kt} \;+\; \varepsilon_{it}. \tag{24}$$

A monotone decline of β_k in k is the testable dose-response prediction. Table 17 and Figure 9 report the results. The pattern is clean. The second-lowest quintile (Q_2) is statistically indistinguishable from the reference across all three outcomes, while Q_3–Q_5 produce monotonically larger declines in cash, sales, and assets. For cash, the gradient runs from -3.4 at Q_2 to -18.6 at Q_5; for assets, from $+11.6$ at Q_2 to -117.2 at Q_5. The absence of an effect at the lowest exposure level rules out the concern that the continuous slope is driven by a common pattern unrelated to real-estate exposure.

Binary high-vs-low splits. As a further check, we re-estimate the difference-in-differences with binary high-vs-low treatment indicators based on three cutoffs of pre-policy exposure: top vs. bottom quartile, top vs. bottom tercile, and above vs. below median (Table 18). Cash and assets are robustly negative across all three cutoffs, with statistical significance strongest in the tercile and median specifications. The dynamic counterpart—a high-vs-low event study on cash holdings (Figure 10)—shows pre-policy coefficients indistinguishable from zero, an immediate post-policy decline, and a gap that deepens to about -12 hundred-million RMB ten quarters out before a

Table 17: Quintile Dose-Response: Post Effect Relative to Q_1

	Cash (100M RMB)	**Sales** (100M RMB)	**Assets** (100M RMB)
$Q_2 \times$ Post	-3.38	-3.04	$+11.61$
	(3.93)	(5.76)	(29.82)
$Q_3 \times$ Post	-12.66^{***}	-10.18^{**}	-54.37^{***}
	(4.69)	(4.11)	(19.77)
$Q_4 \times$ Post	-16.20^{***}	-9.90^{**}	-88.24^{***}
	(4.72)	(4.74)	(25.66)
$Q_5 \times$ Post	-18.58^{***}	-15.27^{**}	-117.21^{***}
	(5.45)	(6.05)	(36.46)
Firm FE	Yes	Yes	Yes
Industry $\times$ Time FE	Yes	Yes	Yes
Observations	142,024	139,429	142,024

Notes. Firms are sorted into quintiles of pre-policy treatment intensity (firm-level, time-invariant). The reference category is Q_1 (lowest exposure). Standard errors clustered at the industry level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

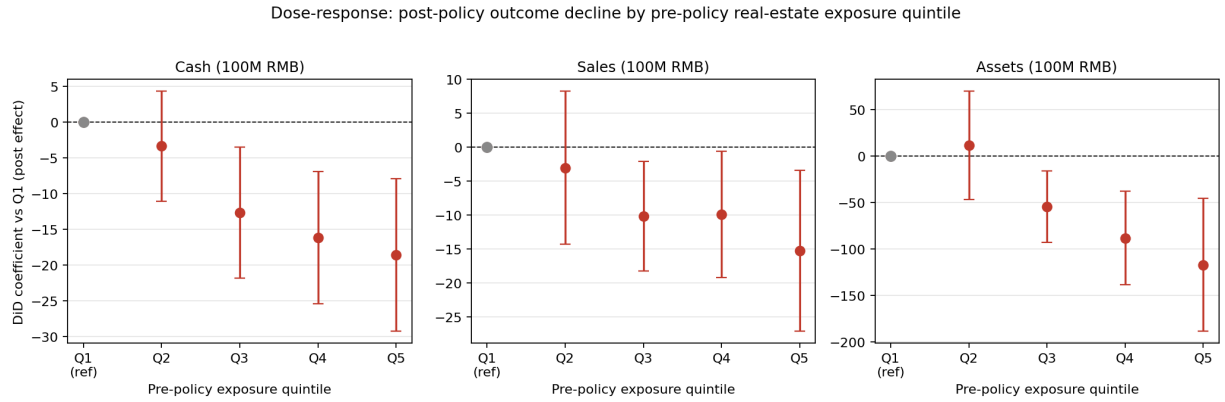


Figure 9: Dose-response by pre-policy real-estate exposure quintile

Notes: Each panel plots the quintile-specific post-policy effect (β_k in (24)) and its 95% confidence interval, with Q_1 (lowest exposure) as the reference. The monotone gradient across cash, sales, and assets and the near-zero effect at Q_2 support a dose-response reading of the continuous-exposure coefficient.

partial recovery.

Interpretation. The continuous-exposure coefficient β in our main specification is therefore best read as a parsimonious summary of how post-policy outcomes vary with pre-determined exposure, rather than as a structural marginal response to an infinitesimal change in exposure. The dose-response gradient and the binary splits together establish the underlying fact—that

Table 18: Binary High-vs-Low DiD: Robustness Across Cutoffs

	Cash (100M RMB)	**Sales** (100M RMB)	**Assets** (100M RMB)
Panel A. Top vs. bottom quartile			
High × Post	-11.37^{*}	-1.71	-47.97^{*}
	(5.97)	(4.87)	(25.32)
Observations	71,375	70,074	71,375
Panel B. Top vs. bottom tercile			
High × Post	-16.07^{***}	-13.63^{*}	-116.30^{**}
	(5.61)	(7.18)	(44.61)
Observations	94,704	92,950	94,704
Panel C. Above vs. below median (full sample)			
High × Post	-10.81^{***}	-8.57^{*}	-88.07^{**}
	(4.10)	(4.97)	(38.81)
Observations	142,024	139,429	142,024
Firm FE	Yes	Yes	Yes
Industry × Time FE	Yes	Yes	Yes

Notes. High = 1 if firm's pre-policy treatment intensity falls in the indicated upper portion of the distribution. Panels A and B compare the top group to the bottom group only (middle group dropped); Panel C uses the full sample. Standard errors clustered at the industry level in parentheses. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

firms more exposed ex ante to real estate through trade-credit channels deteriorate more after the policy—without relying on the stronger assumptions required to interpret the continuous slope causally.

A.4 Trade-Credit Dynamics

If the channel operates through payment delays from real-estate buyers to upstream suppliers, post-policy balance sheets at high-intensity firms should display a characteristic signature: rising accounts-receivable balances, elevated AR-to-sales ratios, and lengthening days-sales-outstanding. Table 19 reports five direct trade-credit metricsAR balance, ln(AR), AR/Sales, DSO measured in days, and a broad trade-credit-to-sales ratio incorporating bills receivable and prepayments.

Each metric shifts in the direction the channel predicts and to a quantitatively large degree. Days-sales-outstanding rises by roughly five days per percentage point of pre-policy intensitya sizeable lengthening relative to the cross-sectional averageand AR-to-sales rises by over five percentage points. The AR balance in level terms also rises significantly. These results provide direct

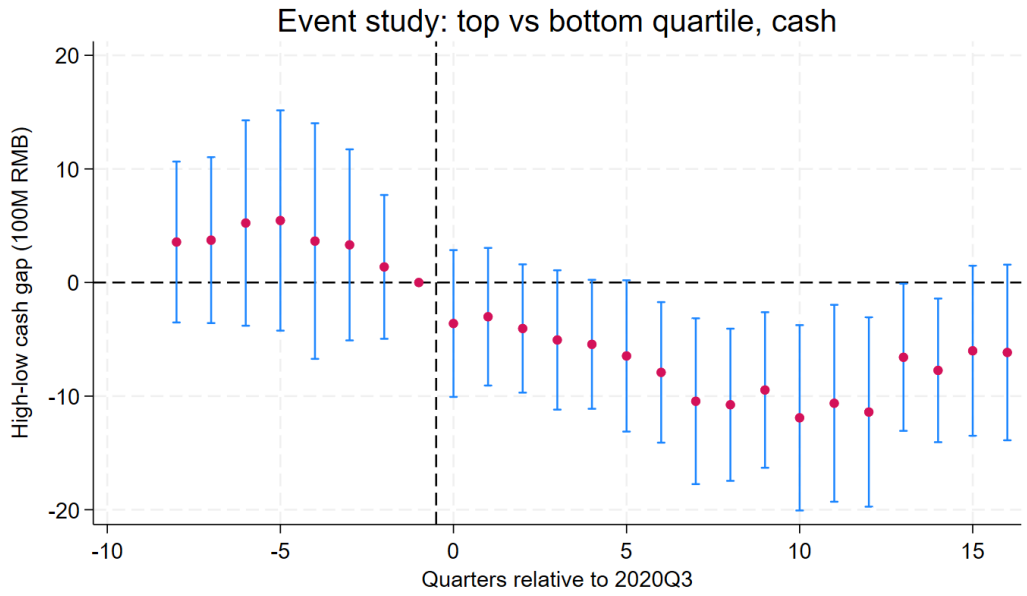


Figure 10: Event study: cash holdings, high vs. low exposure (median cutoff)

Notes: Event-time coefficients from a high-versus-low specification with the median pre-policy exposure as cutoff. Pre-policy coefficients are statistically indistinguishable from zero; the post-policy gap deepens through approximately quarter +10 before partial recovery. 95% confidence intervals shown.

Table 19: Trade-Credit Dynamics

	(1) AR (100M RMB)	(2) ln(AR)	(3) AR/Sales Ratio	(4) DSO Days	(5) TCS Ratio
Intensity(%) × Post	1.7993**	0.0067	0.0544***	4.8943***	0.0720***
	(0.6906)	(0.0119)	(0.0118)	(1.0648)	(0.0134)
N	142024	141288	139419	139419	139419
Firm FE	Yes	Yes	Yes	Yes	Yes
Industry-by-Time FE	Yes	Yes	Yes	Yes	Yes
Adj. R^2	0.873	0.883	0.691	0.691	0.662

Standard errors clustered at industry level in parentheses.

DSO = AR / (quarterly sales / 90). TCS = (AR + BR + Prepayments) / Sales.

Ratios winsorized at 1/99. Firm + Time + Industry-by-Time FE.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

evidence that the channel operates through the payment-delay margin invoked in the theoretical framework, rather than through alternative channels such as demand shocks or price adjustments.

A.5 Real-Economy Outcomes

Beyond financial balance sheets, the channel propagates to real production decisions (Table 20). We examine quarterly capital expenditure (CAPEX), R&D outlay, inventory stock, fixed-asset stock, and gross profit, all in 100M RMB and drawn from the cash-flow, income, and balance-sheet statements.

Table 20: Real-Economy Outcomes

	Dependent Variable (100M RMB)				
	(1) CAPEX	(2) R&D	(3) Inventory	(4) Fixed Assets	(5) Gross Profit
Intensity(%) × Post	-0.08***	-0.05	-11.29***	-0.76	-0.18***
	(0.03)	(0.03)	(1.95)	(0.77)	(0.06)
N	124807	94311	125916	127426	125164
Firm FE	Yes	Yes	Yes	Yes	Yes
Industry-by-Time FE	Yes	Yes	Yes	Yes	Yes
Adj. R^2	0.848	0.790	0.944	0.962	0.931

Standard errors clustered at industry level in parentheses.
CAPEX, R&D, Gross Profit are quarterly flows; Inventory, Fixed Assets are quarter-end stocks.
Firm + Time + Industry-by-Time FE.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

High-intensity firms cut quarterly CAPEX, run down inventory by large amounts, and earn substantially less gross profit. R&D and fixed-asset levels do not respond significantly, consistent with the longer adjustment horizon of these decisionsR&D programs and physical capital adjust slowly relative to working-capital items. The pattern reinforces the channel's interpretation as financial in origin but real in consequence: the trade-credit shock at high-intensity firms reduces working capital, which constrains investment, forces inventory drawdown, and erodes operating profitability.

A.6 Heterogeneity by Financial Constraint

Section 3.6 establishes that financial constraint amplifies the channel within the exposed group of firms. This appendix reports the full heterogeneity analysesacross the complete outcome set and using both ownership and firm size as constraint proxieson which that argument rests.

For ownership, the channel concentrates almost entirely in private firms. Private firms experi-

Table 21: Heterogeneity by Ownership

	Log Outcomes			Balance-Sheet Ratios	
	(1) ln(Cash)	(2) ln(Sales)	(3) ln(Assets)	(4) Cash/Assets	(5) Debt/Assets
Intensity(%) × Post	0.0055	0.0099	0.0066	-0.0001	-0.0011
	(0.0106)	(0.0190)	(0.0163)	(0.0012)	(0.0010)
Intensity(%) × Post × Private	-0.0377***	-0.0236	-0.0258	-0.0021	0.0037***
	(0.0082)	(0.0178)	(0.0156)	(0.0016)	(0.0011)
N	103960	103066	103965	103965	103965
Firm FE	Yes	Yes	Yes	Yes	Yes
Industry × Time FE	Yes	Yes	Yes	Yes	Yes
Adj. R^2	0.860	0.909	0.949	0.618	0.774

Standard errors clustered at industry level in parentheses.

Sample restricted to private and SOE firms. Baseline group = SOE.

Ownership measured 2020-01-01 (pre-policy). Firm + Time + Industry-by-Time FE.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 22: Heterogeneity by Firm Size

	Log Outcomes			Balance-Sheet Ratios	
	(1) ln(Cash)	(2) ln(Sales)	(3) ln(Assets)	(4) Cash/Assets	(5) Debt/Assets
Intensity(%) × Post	-0.0190*	-0.0054	-0.0082	-0.0019***	0.0025***
	(0.0104)	(0.0122)	(0.0083)	(0.0005)	(0.0008)
Intensity(%) × Post × Small/Med	0.0067	0.0062	-0.0138	0.0036**	-0.0069***
	(0.0192)	(0.0232)	(0.0188)	(0.0014)	(0.0011)
N	118295	117163	118301	118301	118301
Firm FE	Yes	Yes	Yes	Yes	Yes
Industry × Time FE	Yes	Yes	Yes	Yes	Yes
Adj. R^2	0.861	0.907	0.949	0.623	0.773

Standard errors clustered at industry level in parentheses.

Sample stratified by Wind size category (2020-01-01). Baseline = Large.

Firm + Time + Industry-by-Time FE.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

ence a substantial decline in log-cash per intensity unitroughly four percent in proportional terms while the corresponding effect for SOEs is statistically indistinguishable from zero. Private firms also raise their debt-to-asset ratio, reflecting an active substitution toward external borrowing to bridge the cash drawdown. SOEs exhibit neither response: with implicit guarantees and preferential bank access, they neither lose cash on the balance sheet nor need to substitute toward debt to smooth the shock.

For firm size, the picture is qualitatively distinct from the ownership comparison. Both large and small firms respond, but through different margins. Large firms substitute toward debt and draw down cash; small and medium firms instead deleverage and contract their balance sheets. The latter pattern reflects an inability to access external borrowing rather than a smaller underlying shock: smaller firms forced into deleveraging are absorbing the same channel through a more contractionary adjustment.

Together, the ownership and size analyses establish that the channel reaches financially constrained firms broadly, with constraint acting as an amplifier rather than the channel itself.

A.7 Network Decomposition: Full Results

Section 3.6 presents the joint network-by-constraint test on the cash outcome. This appendix reports the underlying network decomposition results across additional outcomes and decomposition orders. Table 23 reports the network decomposition on cash with finer-grained components: direct (A), second-order (A^2), and the cumulative higher-order residual. Table 24 repeats the analysis on sales. Table 25 reports the ownership-by-network interaction with sales as the outcome.

Across outcomes the dominance of higher-order over direct exposure is robust. The isolated second-order term A^2 is imprecisely estimated because of collinearity with the residual higher-order terms; pooling all multi-step paths into a single cumulative component delivers the sharpest estimates. Sales as an outcome reproduces the same qualitative pattern as cashdirect exposure negligible, higher-order exposure substantialwith the joint ownership interaction showing the same constraint asymmetry.

Table 23: Network Decomposition Cash Outcome

	Cash (100M RMB)				
	(1) Total	(2) Direct	(3) 2nd order	(4) Higher	(5) Direct+Higher
Intensity[total] × Post	-2.362***				
	(0.536)				
Intensity[direct] × Post		-1.785			-0.063
		(1.693)			(0.119)
Intensity[2nd order] × Post			-6.301		
			(4.868)		
Intensity[higher order] × Post				-2.948***	-2.934***
				(0.294)	(0.301)
N	142024	142024	142024	142024	142024
Adj. R^2	0.899	0.899	0.899	0.899	0.899

Intensity = AR/Assets × Leontief component. Direct = A; 2nd = A^2; Higher = A^2+A^3+...

Firm + Time + Industry-by-Time FE. Cluster SE at industry.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 24: Network Decomposition Sales Outcome

	Sales (100M RMB)			
	(1) Total	(2) Direct	(3) Higher	(4) Direct+Higher
Intensity[total] × Post	-2.037**			
	(0.800)			
Intensity[direct] × Post		-1.336		0.104
		(1.500)		(0.225)
Intensity[higher] × Post			-2.474***	-2.498***
			(0.257)	(0.252)
N	139429	139429	139429	139429
Firm FE	Yes	Yes	Yes	Yes
Industry-by-Time FE	Yes	Yes	Yes	Yes

Standard errors clustered at industry level in parentheses.

Intensity = AR/Assets × Leontief component. Direct = A; Higher = A^2+A^3+$\cdots$.

Firm + Time + Industry-by-Time FE.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 25: Joint Network and Ownership Sales Outcome

	Sales (100M RMB)				
	(1) Higher	(2) Higher×Priv	(3) Direct	(4) Direct×Priv	(5) Both
Higher × Post	-2.474***	-2.920***			-3.740***
	(0.257)	(0.710)			(0.687)
Higher × Post × Private		1.387			2.183**
		(1.072)			(0.904)
Direct × Post			-1.336	0.850	3.828***
			(1.500)	(3.256)	(0.941)
Direct × Post × Private				-1.716	-3.639***
				(2.561)	(1.298)
N	139429	103069	139429	103069	103069
Firm FE	Yes	Yes	Yes	Yes	Yes
Industry-by-Time FE	Yes	Yes	Yes	Yes	Yes

Standard errors clustered at industry level in parentheses.
Private firms = (Wind 2020-01-01 snapshot). Baseline = SOE.
Cols. 2,4,5 restrict to SOE+Private (other categories dropped).
Firm + Time + Industry-by-Time FE.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.8 Dyadic Spillover and Identification Choice

The main empirical strategy uses an ecological treatment intensity: firm-level pre-policy AR ratios interacted with industry-level Leontief exposure to real estate. An alternative would use direct dyadic exposurematching each supplier to its specific top-five disclosed AR debtors, identifying which debtors are RE firms, and weighting by the disclosed share. Table 26 implements this dyadic approach using Wind's top-five AR-debtor disclosure (5,456 firms, 2018–2020 pre-policy snapshot); Table 27 compares the two measures head-to-head.

The dyadic measure does not displace the ecological one. When entered jointly, the ecological coefficient retains its magnitude and significance while the dyadic coefficient remains indistinguishable from zero. Two factors explain the gap. First, Wind's top-five disclosure captures only the tail of AR exposure for most listed firms, so the dyadic measure has substantial measurement error in the unobserved buyer relationships. Second, the dyadic measure is sparse: only a small fraction of firms have an identified RE customer in the top five at all, and a smaller fraction still have a customer whose breach status can be matched to ownership data. The ecological measure

Table 26: Dyadic Spillover via Wind Top-Five AR Debtors

	Dependent Variable (100M RMB)				
	(1) Cash	(2) Sales	(3) Assets	(4) Debt	(5) Equity
RE-AR Share(%) × Post	0.117	-0.298	-2.931	-0.635	-1.032**
	(0.496)	(0.377)	(2.532)	(0.676)	(0.393)
N	142024	139429	142024	142024	142024
Firm FE	Yes	Yes	Yes	Yes	Yes
Industry-by-Time FE	Yes	Yes	Yes	Yes	Yes
Adj. R^2	0.899	0.951	0.953	0.897	0.976

Standard errors clustered at industry level.

Treatment = pre-policy (2018-2020) percentage of top-5 AR owed by RE-sector debtors. Source: Wind.

Firm + Time + Industry-by-Time FE.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 27: Ecological vs. Dyadic Head-to-Head

	Cash (100M RMB)			Sales (100M RMB)		
	(1) Ecological	(2) Dyadic	(3) Both	(4) Ecological	(5) Dyadic	(6) Both
Intensity(%) × Post [ecological]	-2.362***		-2.363***	-2.037**		-2.035**
	(0.536)		(0.535)	(0.800)		(0.800)
RE-AR Share(%) × Post [dyadic]		0.117	0.119		-0.298	-0.297
		(0.496)	(0.500)		(0.377)	(0.379)
N	142024	142024	142024	139429	139429	139429
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Industry-by-Time FE	Yes	Yes	Yes	Yes	Yes	Yes

Standard errors clustered at industry level in parentheses.

Ecological intensity = pre-policy AR/Assets × Leontief exposure to K70.

Dyadic intensity = pre-policy share of top-5 AR owed by RE debtors (Wind, 2018–2020).

Firm + Time + Industry-by-Time FE.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

therefore captures the channel more cleanly than the disclosed dyadic alternative, which validates rather than displaces the main paper's identification strategy.

B Proof of Proposition 1

Throughout, both borrowing constraints bind and prices $(W, P, P_h, R_h, R_n, \theta_n)$ and technology (A_h, A_n, α) are taken as given. We prove parts (i)–(iv) in turn.

Part (i): the prepayment first-order condition. The downstream firm solves

$$\max_{X,\phi} \; P_h A_h X^{\alpha} - PX - G(\phi)PX - R_h \phi PX \quad \text{s.t.} \quad \phi PX \leq \theta_h P_h A_h X^{\alpha},$$

with Lagrangian

$$\mathcal{L} = P_h A_h X^{\alpha} - PX - G(\phi)PX - R_h \phi PX + \lambda_h [\theta_h P_h A_h X^{\alpha} - \phi PX] .$$

The first-order condition with respect to ϕ is $\partial \mathcal{L}/\partial \phi = -G'(\phi)PX - R_h PX - \lambda_h PX = 0$, i.e. $-G'(\phi) = R_h + \lambda_h$. With $G(\phi) = \frac{1}{2}(1-\phi)^2$ we have $G'(\phi) = -(1-\phi)$, hence $1-\phi = R_h + \lambda_h$ and $\phi = 1 - R_h - \lambda_h$. Because the constraint binds, $\lambda_h > 0$. This proves (i).

Reduction to a single equation in ϕ. The first-order condition with respect to X is

$$\alpha P_h A_h X^{\alpha-1} (1 + \lambda_h \theta_h) = P \left[1 + G(\phi) + \phi \left(R_h + \lambda_h\right)\right]. \tag{1}$$

Using $R_h + \lambda_h = 1 - \phi$ and $G(\phi) = \frac{1}{2}(1-\phi)^2$, the bracket on the right simplifies:

$$1 + \tfrac{1}{2}(1-\phi)^2 + \phi(1-\phi) = 1 + (1-\phi)\left[\tfrac{1}{2}(1-\phi) + \phi\right] = 1 + \tfrac{1}{2}(1-\phi^2) = \tfrac{3-\phi^2}{2}.$$

The binding constraint $\phi PX = \theta_h P_h A_h X^{\alpha}$ gives $P_h A_h X^{\alpha-1} = \phi P/\theta_h$. Substituting both into (1) and cancelling P,

$$\alpha \, \frac{\phi}{\theta_h} \, (1 + \lambda_h \theta_h) = \frac{3 - \phi^2}{2}.$$

Finally substituting $\lambda_h = 1 - R_h - \phi$ from part (i) and rearranging yields a single implicit equation in ϕ and θ_h:

$$F(\phi, \theta_h) \;\equiv\; \frac{\alpha\phi}{\theta_h} + \alpha\phi\,(1 - R_h) + \phi^2\left(\tfrac{1}{2} - \alpha\right) - \tfrac{3}{2} \;=\; 0. \tag{2}$$

Part (ii): comparative statics. Differentiating (2) implicitly,

$$\frac{d\phi}{d\theta_h} = -\frac{\partial F/\partial\theta_h}{\partial F/\partial\phi} = \frac{\alpha\phi/\theta_h^2}{\alpha/\theta_h + \alpha(1 - R_h) + \phi\,(1 - 2\alpha)}.$$

The numerator is strictly positive. The denominator is strictly positive whenever $\alpha \le \frac{1}{2}$ (all three terms are then non-negative and the first is positive); for $\alpha > \frac{1}{2}$ it remains positive provided $\alpha/\theta_h + \alpha(1 - R_h) > (2\alpha - 1)\phi$, which holds in the binding regime because $\theta_h < 1$ makes α/θ_h large relative to $(2\alpha - 1)\phi < 1$. Hence $d\phi/d\theta_h > 0$. By part (i), $\lambda_h = 1 - R_h - \phi$, so $d\lambda_h/d\theta_h = -\,d\phi/d\theta_h < 0$. This proves (ii).

Numerical illustration. With $\alpha = 0.5,\ R_h = 0.02,\ P = P_h A_h = 1$, equation (2) gives, in the binding region, $\phi = 0.696,\ \lambda_h = 0.285$ at $\theta_h = 0.30$ and $\phi = 0.602,\ \lambda_h = 0.378$ at $\theta_h = 0.25$: tightening θ_h raises λ_h and lowers ϕ, as claimed.

Part (iii): transmission. From (10), $Y_n = \phi PX/(W/A_n - \theta_n P)$ with $W/A_n - \theta_n P > 0$ in the constrained regime, so $\partial Y_n/\partial\phi = PX/(W/A_n - \theta_n P) > 0$. Holding the downstream order X fixed, the chain rule gives

$$\left.\frac{\partial Y_n}{\partial\theta_h}\right|_{\text{prepayment channel}} = \frac{\partial Y_n}{\partial\phi}\cdot\frac{d\phi}{d\theta_h} > 0,$$

using $d\phi/d\theta_h > 0$ from part (ii). Thus a fall in θ_h lowers ϕ and, through the prepayment, lowers Y_n. This proves (iii).

Part (iv): exogenous trade credit. If ϕ is fixed at a constant $\bar{\phi}$, it no longer solves (8); $d\phi/d\theta_h = 0$ by construction. Then $\partial Y_n/\partial\phi \cdot d\phi/d\theta_h = 0$, so the prepayment channel of part (iii) carries no transmission, and the borrowing-constraint shock affects the upstream sector only through any change in the downstream's physical order X—the conventional demand channel. This proves (iv).

C Seven-Sector Aggregation Map

This appendix documents the mapping from the National Bureau of Statistics industry-classification codes (which are inherited by both the 2020 China input-output table and the CSMAR firm-level financial database) to the seven aggregated sectors used in the multi-sector model of Section 4.2. The aggregation is designed to align the IO-table-based calibration of Section 4.3 with the empirical analysis of Section 3, where the same seven-sector grouping is used to construct the Leontief-exposure component of the treatment intensity measure.

Table 28 reports the mapping. Three remarks are worth making.

First, the manufacturing letter-code C is split into Basic Manufacturing (C13–C31) and Advanced Manufacturing (C32–C43). The split follows the conventional division between material-processing and equipment-manufacturing branches: C13–C31 covers material-processing sectors (food, textile, wood, paper, petroleum, chemical, rubber, non-metal mineral products), while C32–C43 covers equipment-producing sectors (metal products, machinery, electronics, transport equipment, electrical equipment). The two-way split is empirically meaningful in our context because the two groups have very different production-network positions: Basic Manufacturing is heavily reliant on intermediate inputs from other commodity sectors (input share $\approx$ 75%), while Advanced Manufacturing is more value-added intensive and more downstream.

Second, only the K70 sub-code is mapped to Real Estate. The full K letter-code spans K70 (real estate), K71 (leasing of real estate), and K72 (real estate intermediary services); we restrict to K70 because our model's real-estate sector represents new-housing development, which corresponds in the IO classification to K70 only. K71 and K72 are aggregated into the Services sector together with the other service-industry codes.

Third, the financial industry letter-code J is excluded from the seven-sector real-economy model. Section 3 likewise excludes J-coded firms from the CSMAR estimation sample.

D Intermediate Input Share Matrix

This appendix reports the 7×7 matrix of intermediate input shares ζ_{ij} used in the multi-sector model of Section 4.2, where row i lists the intermediate-input bundle of sector i: ζ_{ij} is the share of

Table 28: Seven-sector aggregation of NBS industry codes

#	Sector	NBS letter / digit codes
1	Agriculture	A (A01–A05): agriculture, forestry, animal husbandry, fishery
2	Mining	B (B06–B12): all mining and quarrying sub-sectors
3	Basic Manufacturing	C13–C33: food, textiles, wood, paper, petroleum, chemicals, rubber, non-metallic minerals
4	Adv. Manufacturing	C34–C40: metal products, machinery, electronics, transport equipment, electrical equipment
5	Construction	E (E47–E50): construction of buildings, civil engineering, installation, decoration
6	Real Estate	K70: real estate development
7	Services	D, F, G, H, I, K71–K72, L, M, N, O, P, Q, R: utilities, wholesale and retail, transport and post, hospitality, IT services, real-estate leasing/intermediary, leasing and business services, scientific R&D, environment, residential services, education, health, culture

Notes. NBS = National Bureau of Statistics. The classification follows the GB/T 4754–2017 standard, which is the common classification used by both the 2020 China Input-Output Table and CSMAR. The mapping is identical to that used in Section 3 for the construction of firm-level treatment intensity.

sector j's output in sector i's intermediate input aggregator $X_{i,t} = \prod_j (X^C_{ij,t}/\zeta_{ij})^{\zeta_{ij}}$. Two modifications are imposed on the raw shares computed from the 2020 input-output table.

First, the real-estate column is set to zero ($\zeta_{i,h} = 0$ for every i), reflecting the model's treatment of new housing as a final-use durable asset rather than a productive intermediate good. The non-trivial entries that the raw IO table records for sector-to-RE flows (largest in the RE-to-RE diagonal at 0.149) primarily reflect commercial leasing services and imputed rentals on owner-occupied housing, not productive use of new construction as an intermediate input; we treat these as conceptually distinct from the model's ζ matrix and reassign them to the Services sector through the appropriate IO aggregation.

Second, each row is renormalized so that $\sum_{j \neq h, j \neq i} \zeta_{ij} = 1$. The own-sector entry ζ_{ii} is also set to zero before renormalization, consistent with the model specification that sector i's intermediate aggregator runs over $j \neq i$. Real estate's own input share vector $\zeta_{h,i}$, $i \neq h$, is also renormalized after the within-RE entry is zeroed.

Table 29 reports the resulting matrix. Each row sums to one (entries strictly less than one due to the column-h and diagonal zeroing). Empty cells correspond to structurally zero entries (ζ_{ii} and

Table 29: Intermediate input share matrix ζ_{ij} (7×7, row-normalized; diagonal and real-estate column are zero by construction)

	Supplier sector j						
Buyer i	Agri	Min	BasicM	AdvM	Constr	RE	Serv
Agriculture	—	0.001	0.406	0.027	0.002	0	0.205
Mining	0.001	—	0.247	0.102	0.001	0	0.374
Basic Manuf.	0.124	0.111	—	0.020	0.000	0	0.215
Adv. Manuf.	0.000	0.001	0.261	—	0.000	0	0.202
Construction	0.009	0.012	0.515	0.069	—	0	0.347
Real Estate	0.001	0.000	0.069	0.003	0.056	—	0.870
Services	0.015	0.027	0.197	0.095	0.006	0	—

Notes. Entries are intermediate input shares from the 2020 China Input-Output Table, aggregated to the seven sectors of Appendix C, with two transformations applied: (i) the real-estate column ($j = h$) is set to zero, reflecting the model's treatment of new housing as a final-use asset; (ii) each row is renormalized so that $\sum_{j \neq h, j \neq i} \zeta_{ij} = 1$. Row sums for non-RE buyers therefore equal one; the renormalization absorbs the small ($< 3\%$ for any non-RE row) mass originally assigned to the RE column. For real estate, the within-RE entry (14.9% in the raw IO table) is reassigned to the other six sectors proportionally, resulting in the row dominated by Services (0.87).

$\zeta_{i,h}$).

A useful object for the empirical analysis is the column-h slice of the Leontief inverse, $\Psi_{i,h} = [(I - \Omega)^{-1}]_{i,h}$, where $\Omega_{ij} = (1 - \alpha_i)\zeta_{ij}$. This entry measures the total (direct plus indirect) increase in sector i's output required to produce one unit of additional real-estate final demand. It is the network-exposure component of the treatment intensity used in Section 3.

E Derivation of the Empirical Implied Loss Ratio

This appendix derives equation (23) of Section 5.1 from primitive accounting identities, and discusses the first-order approximation underlying the comparison between EILR and the model's cumulative percentage GDP deviation.

Notation

For each post-policy quarter $t \in \{1, \ldots, Q_{\text{post}}\}$ and each listed firm $i \in \{1, \ldots, N\}$ in the Table 2 sample, let:

- $I_i \geq 0$ denote the firm's treatment intensity in percentage-points, defined as in Section 3 by $I_i = 100 \cdot (\text{AR}/\text{Assets})_i^{\text{pre}} \cdot \Psi_{i,h}$, where $\Psi_{i,h}$ is the Leontief exposure to real estate. By construction I_i is firm-specific and time-invariant.

- β_S be the post $\times$ intensity coefficient from column (2) of Table 2, with units of 100M per firm-quarter per intensity percentage-point.
- S_{it} be firm i's realized quarterly sales in 100M; $\bar{S} \equiv (Q_{\text{post}} N)^{-1} \sum_{i,t} S_{it}$ the post-period sample mean.
- $\bar{I} \equiv N^{-1} \sum_i I_i$ the cross-sectional mean of intensity.
- GDP_t be nominal GDP in quarter t, and $\lambda \equiv \left(\sum_{i,t} S_{it}\right) / \left(\sum_t \text{GDP}_t\right)$ the share of GDP accounted for by the listed-firm sample over the post-policy window.

From per-firm marginal effect to aggregate sales loss

The reduced-form specification in Section 3 imposes that, for any firm i in any post-policy quarter t, the post $\times$ intensity coefficient gives the marginal effect on sales:

$$\mathbb{E}\left[\Delta S_{it} \,\middle|\, i \in \text{sample},\, t \in \text{post}\right] = \beta_S \cdot I_i. \tag{3}$$

Here ΔS_{it} denotes the policy-induced deviation of firm i's quarter-t sales from its counterfactual (no-policy) value.

Summing (3) across the N firms in the sample at any post-policy quarter t:

$$\Delta S_t^{\text{listed}} \;=\; \sum_{i=1}^{N} \beta_S \, I_i \;=\; \beta_S \cdot N \cdot \bar{I}. \tag{4}$$

Equation (4) is the per-quarter aggregate sales loss across the listed-firm sample, in 100M.

From listed-firm sales to GDP

Total listed-firm sales in any post-policy quarter t are, by definition,

$$S_t^{\text{listed}} \;\equiv\; \sum_{i=1}^{N} S_{it} \;=\; N \cdot \bar{S}_t, \tag{5}$$

where $\bar{S}_t \equiv N^{-1} \sum_i S_{it}$ is the cross-sectional mean of quarterly sales in quarter t. Using the time average $\bar{S} = Q_{\text{post}}^{-1} \sum_t \bar{S}_t$ as a representative value, and the definition of λ together with the time-

aggregation identity $\sum_t S_t^{\text{listed}} = \lambda \sum_t \text{GDP}_t$, we obtain the representative-quarter relationship

$$\overline{\text{GDP}} \;=\; \frac{N \cdot \bar{S}}{\lambda}, \tag{6}$$

where $\overline{\text{GDP}} = Q_{\text{post}}^{-1} \sum_t \text{GDP}_t$.

The EILR

Dividing the aggregate sales loss (4) by the representative-quarter GDP (6) and multiplying by 100 to express the result in percentage points yields

$$\text{EILR}_q \;\equiv\; \frac{\Delta S_t^{\text{listed}}}{\overline{\text{GDP}}} \cdot 100 \;=\; \frac{\beta_S \cdot N \cdot \bar{I}}{N \cdot \bar{S}/\lambda} \cdot 100 \;=\; \frac{\beta_S \cdot \bar{I} \cdot \lambda}{\bar{S}} \cdot 100, \tag{7}$$

which is equation (23). The number of firms N cancels exactly because it enters the numerator linearly through cross-firm summation and the denominator linearly through listed-firm aggregate sales. Thus EILR_q depends only on the regression coefficient and four sample moments $(\bar{I}, \bar{S}, \lambda, \beta_S)$, with no dependence on sample size — a desirable invariance property.

The cumulative version over the Q_{post} post-policy quarters is

$$\text{EILR}_{\text{cum}} \;\equiv\; \sum_{t=1}^{Q_{\text{post}}} \text{EILR}_q \;=\; Q_{\text{post}} \cdot \text{EILR}_q, \tag{8}$$

which is comparable to the model's cumulative percentage GDP deviation $\sum_t 100 \cdot (\text{GDP}_t/\text{GDP}^{\text{ss}} - 1)$ over the same window.

Numerical inputs (Three Red Lines sample)

Substituting the values reported in Section 5.1,

$$\beta_S = -2.04, \quad \bar{I} = 0.6925 \text{ pp}, \quad \bar{S} = 33.10 \text{ 100M}, \quad \lambda = 0.4735, \quad Q_{\text{post}} = 20,$$

into (7)–(8) gives $\text{EILR}_q = -2.021\%$ and $\text{EILR}_{\text{cum}} = -40.43$ pp·q.

Sum-of-deviations versus ratio-of-cumulatives

Throughout we compare EILR_{cum} to the model's sum of per-quarter percentage deviations $\sum_t 100\,(\text{GDP}_t/\text{GDP}^{\text{ss}} - 1)$ rather than to the ratio of cumulative losses, $100 \cdot \sum_t(\text{GDP}_t - \text{GDP}^{\text{ss}}) / \sum_t \text{GDP}^{\text{ss}}$. The two coincide when the percentage deviations are constant across t, and differ by second-order terms otherwise. For the Three Red Lines IRF over the 20-quarter window, the levels-based denominator yields -2.31% against the sum-of-deviations value of -2.02% — a discrepancy of less than 0.3 percentage points, which does not affect any of the conclusions of Section 5.1.

Window dependence of the cumulative comparison

The per-quarter ratio EILR_q does not depend on Q_{post}: extending the post-policy window enters $\bar{S}$ and $\bar{I}$ symmetrically, leaving the per-quarter rate unchanged. The *cumulative* comparison, however, is not a pure scaling choice. By (8), EILR_{cum} grows linearly in Q_{post}, because the pooled empirical effect is constant per quarter, whereas the model's cumulative GDP deviation plateaus as output recovers toward steady state. The ratio of the two therefore declines with the window: the baseline model accounts for about 94% of the implied loss at a 16-quarter horizon but about 79% at the full 20-quarter horizon. This window dependence reflects a genuine difference in dynamics — the empirical spillover is more persistent than the model's recovering response — rather than an artifact of the identity. It is the reason we headline the window-invariant channel decomposition (the 40% share and 1.7× amplification) rather than the match level itself.

F Additional Event Studies

This appendix collects four additional event-study figures supporting the empirical patterns documented in Section 3: Figure 11 shows the dynamic effect of the policy on accounts payable in the cross-industry RE-vs-all specification; Figure 12 shows the corresponding dynamic effect on short-term bank loans; Figure 13 shows the within-RE event study comparing breached and green-tier developers; and Figure 14 shows the supply-side event study for disclosed real-estate suppliers.

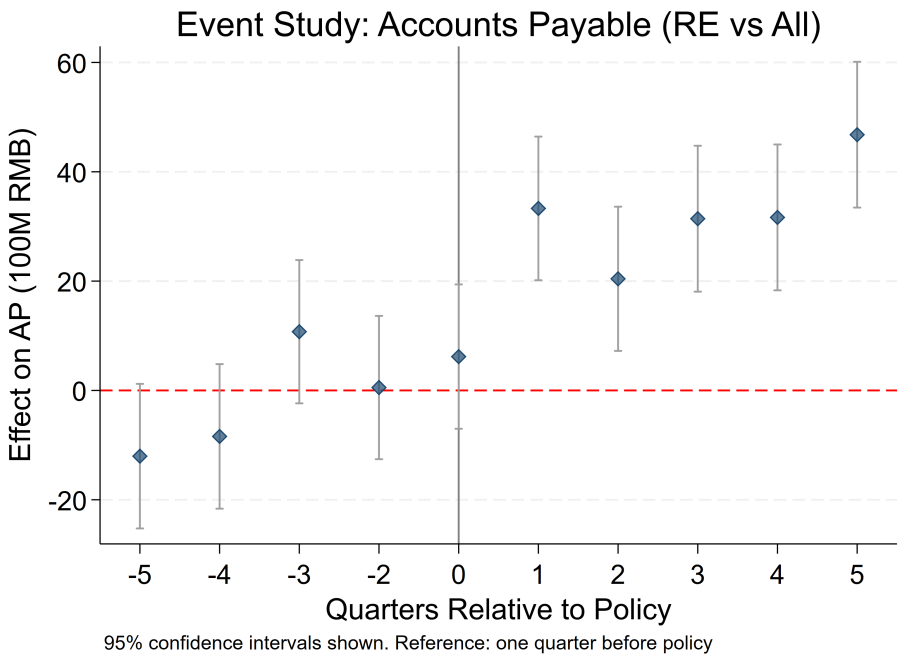


Figure 11: Event Study: Accounts Payable (RE vs. all)

Notes: Coefficients from regressing AP (100M RMB) on RE $\times$ period dummies. Reference: one quarter before policy. Bars represent 95% confidence intervals.

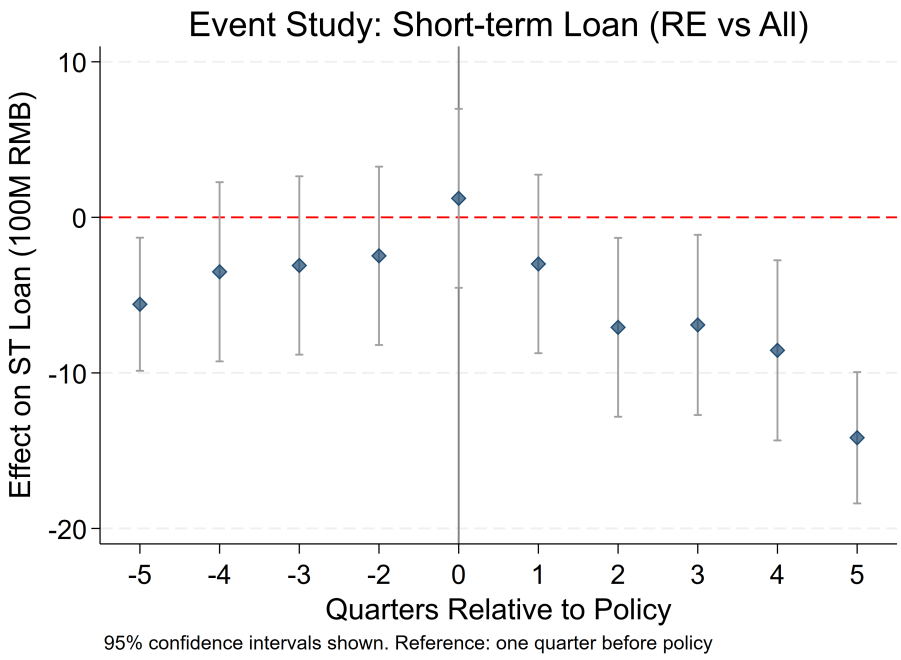


Figure 12: Event Study: Short-term Loan (RE vs. all)

Notes: Coefficients from regressing short-term loans (100M RMB) on RE $\times$ period dummies. Reference: one quarter before policy. Bars represent 95% confidence intervals.

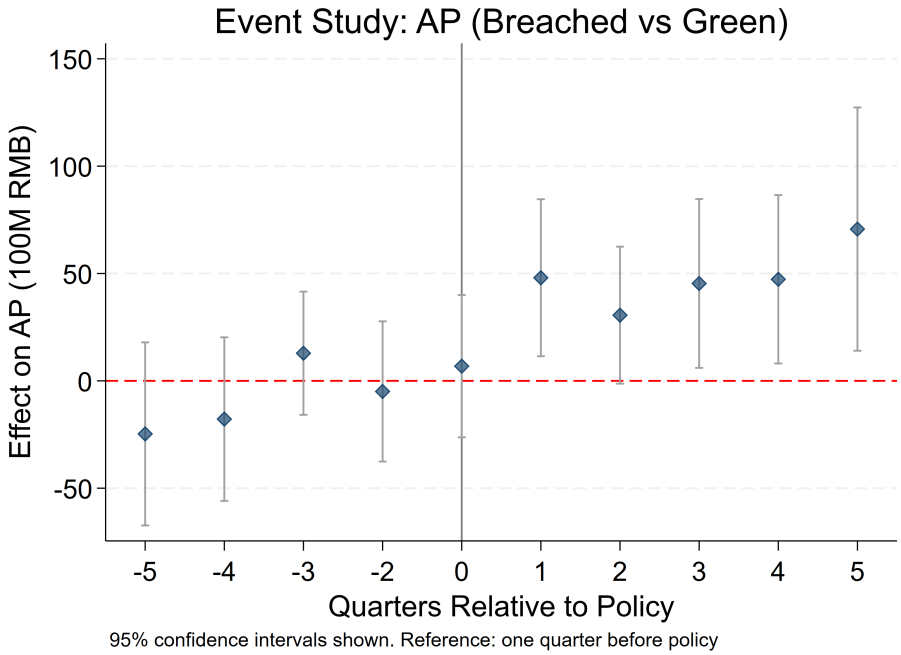


Figure 13: Event Study: AP (Breached vs. green within RE)

Notes: Coefficients from regressing AP (100M RMB) on Breached × period dummies within RE firms only. Reference: one quarter before policy. Bars represent 95% confidence intervals. Standard errors clustered at the firm level.

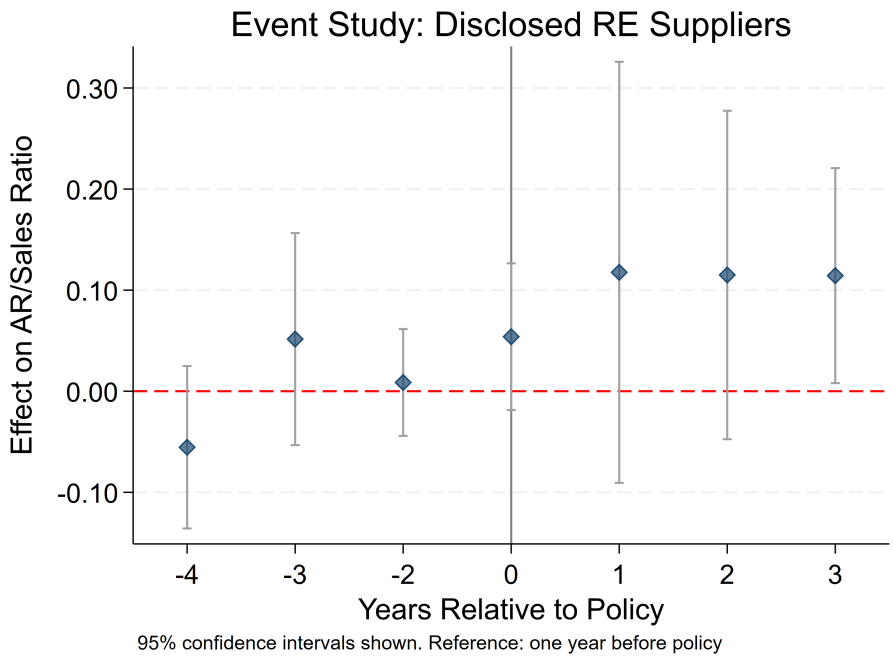


Figure 14: Event Study: AR/Sales (Disclosed RE suppliers)

Notes: Coefficients from regressing AR/Sales on Treat × period dummies. Annual data, balanced panel. Reference: one year before policy (2020). Bars represent 95% confidence intervals.