Trade Credit Transmission of Financial Shocks: Theory and Evidence from China's Real Estate Deleveraging*

Haoxuan Su† Anxu Wang‡ Yang Yao§

June 1, 2026

Abstract

We examine how financial shocks propagate through production networks via trade credit—a major but understudied transmission channel. We exploit China's 2021 "Three Red Lines" policy, which imposed leverage constraints on real estate developers, as a quasi-natural experiment. Combining firm-level financial data with input-output network structure, we document significant and persistent spillover effects: firms more exposed to real estate through trade credit channels experienced declines in cash holdings and sales that persist for over three years. Evidence from both sides of the trade credit relationship—increased payment delays by constrained real estate firms and rising receivables at their suppliers—identifies trade credit as the operative transmission channel. To formalize this mechanism, we develop a tractable production-network model in which the trade-credit terms on each link are set by firms' binding borrowing constraints—so trade credit is the by-product of financial frictions on the network's own edges, the very object that carries a financial shock upstream. The calibrated model generates a GDP trough of -4.7% with output losses persisting over four years, driven primarily by upstream propagation through trade credit. It reproduces most of the empirically implied loss, amplifying the response by about 70% relative to a model with fixed trade credit. The findings imply that sector-specific financial regulation must be evaluated for its network spillovers, not merely its effect on targeted firms; and the framework is portable to other settings where a sector-specific financial shock meets a production network.

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

*We thank seminar participants at the National School of Development, Peking 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

Modern economies are bound together by dense networks of production relationships, financial linkages, and credit exposures. A growing body of research shows that these interconnections can substantially amplify sector-specific shocks: productivity disruptions in one industry cascade downstream through input-output linkages (Long and Plosser, 1983; Acemoglu et al., 2012; Carvalho, 2014), demand shocks to central sectors generate economy-wide multiplier effects (Baqaee and Farhi, 2019), financial distress at key institutions transmits through counterparty exposures (Acemoglu et al., 2015), and bank credit shocks propagate along the production network in both upstream and downstream directions (Huremović et al., 2026). Understanding how shocks propagate through these networks is essential for assessing regulatory interventions and explaining aggregate fluctuations.

Among the channels through which shocks travel across firms, trade credit—the practice of extending or receiving deferred payment between trading partners—stands out as quantitatively important yet empirically understudied. Trade credit accounts for over half of firms' short-term liabilities across OECD countries; in the United States alone, it exceeds bank loans by a factor of three and commercial paper by a factor of fifteen (Altinoglu, 2021). In China, trade credit is similarly an important source of inter-firm financing, with accounts payable accounting for approximately one-fifth of listed firms' total liabilities and an even larger share for non-state-owned firms (Cun et al., 2022). An extensive empirical literature documents that trade credit responds endogenously to shocks to bank credit (Petersen and Rajan, 1997; Cuñat, 2007; Garcia-Appendini and Montoriol-Garriga, 2013) and propagates financial distress along the supply chain (Jacobson and von Schedvin, 2015; Costello, 2020; Almeida et al., 2024). This ubiquity means that when a firm faces financial distress and delays payments to its suppliers, it effectively converts its own liquidity shortage into a credit supply reduction for its trading partners—a transmission mechanism that can propagate distress through supply chains even when physical production continues uninterrupted.

Despite its potential significance, direct empirical evidence on trade credit as a vehicle for shock transmission remains scarce. The core obstacle is data: identifying credit transmission between specific firm pairs requires bilateral information on payment flows that is rarely avail-

able, in China or elsewhere. On the theoretical side, a growing literature integrates *endogenous* trade credit into production-network models—either as supplier-provided credit collateralized by downstream revenue (Altinoglu, 2021), or as a cost-minimizing choice between bank and supplier credit (Reischer, 2020; Luo, 2020; Cun et al., 2022)—but these frameworks are largely static. How endogenous trade credit transmits a *financial* shock that tightens firms' borrowing constraints, dynamically and through the production network, remains an open question.

This paper studies how sector-specific financial shocks transmit through production networks via trade credit, combining novel empirical evidence with a tractable dynamic model that grounds endogenous trade credit in binding financial constraints and applies it quantitatively to the Chinese real-estate context through a multi-sector model. Our empirical setting is China's "Three Red Lines" policy, introduced in August 2020 and implemented starting in 2021, which imposed strict leverage constraints on real estate developers. This policy offers a compelling quasi-natural experiment. The announcement came abruptly with little advance warning, ruling out anticipation effects. Beyond timing, the policy targeted a well-defined sector—real estate development—creating sharp distinctions between treated and untreated firms. Within that treated group, intensity varied across developers based on pre-policy financial metrics, providing dose-response identifying variation. Beyond identification, the affected sector sits at the center of China's production network, with extensive upstream linkages to construction, basic manufacturing, and services—so documented spillovers reflect a macroeconomically meaningful channel rather than an isolated case. Finally, the shock was large enough to generate measurable effects while remaining contained enough to trace transmission through specific supply chains rather than through generalized financial panic.

We begin with an empirical analysis that proceeds in four steps. Our main result establishes the economy-wide reach of the shock. We construct a treatment intensity measure that combines each firm's pre-determined reliance on trade credit (measured by accounts receivable relative to assets before the policy) with its industry's input-output linkage to real estate through the production network. This measure captures a firm's vulnerability to the real estate credit shock transmitted through trade credit channels, while the interaction of firm-level and industry-level variation—together with industry-by-time fixed effects—allows us to isolate the trade credit transmission channel from confounding industry trends. We find substantial spillover effects: a

one-percentage-point increase in exposure is associated with a decline of 236 million RMB in cash holdings and 204 million RMB in sales, with effects emerging sharply after the policy and persisting for over three years—an event-study pattern that is one of the central empirical findings of the paper.

We then provide direct evidence on the transmission mechanism from both sides of the trade credit relationship. On the demand side, we show that the policy shock was real and consequential for targeted firms. Real estate developers experienced severe financial pressure: short-term bank borrowing contracted by 823 million RMB, cash holdings fell by 1.31 billion RMB, while accounts payable surged by 4.58 billion RMB. These patterns indicate that, as formal credit dried up, real estate firms increasingly delayed payments to suppliers, effectively shifting liquidity pressure upstream. Within the real estate sector, firms breaching more regulatory thresholds exhibited proportionally larger effects, confirming that the policy's tiered structure generated dose-response variation. On the supply side, we exploit disclosed customer-supplier relationships from annual reports and find that firms with real estate customers experienced a 9.7 percentage-point increase in accounts receivable relative to sales, alongside declines in sales (log -0.21) and total assets (log -0.11). We corroborate these micro-level findings using industry-level input-output linkages across the full sample of listed firms. The consistency between the demand-side evidence (rising accounts payable at real estate firms) and the supply-side evidence (rising accounts receivable at their suppliers) validates trade credit as the transmission channel: what appears as delayed payment on one firm's balance sheet shows up as involuntary credit extension on another's.

Fourth, we characterize the channel's structure: the spillover travels through higher-order network paths rather than direct sales to real estate, and concentrates among financially constrained private firms rather than state-owned ones. The effect is largest for firms that are both network-exposed and financially constrained. This is what a model of endogenous, constraint-driven trade credit predicts, and is hard to reconcile with models that fix payment terms or ignore firms' financial constraints.

Building on these empirical findings, we develop a tractable dynamic production-network model in which trade credit terms are endogenously determined by firms' financing decisions. When a firm's borrowing constraint tightens, it optimally reduces prepayment to its suppliers, transmitting financial pressure upstream through the network. This endogeneity is precisely what

our empirical evidence supports, and it is what allows the model to replicate the documented payment-delay response. Beyond the application to China, the framework provides a portable tool for studying endogenous trade-credit propagation in any production network economy.

We present the model in two stages. First, we develop a two-sector model that makes the core mechanism analytically transparent: a tightening of the downstream sector's borrowing limit raises its shadow cost of bank credit, induces it to reduce prepayment, and propagates the financial pressure upstream through the borrowing constraint of the supplier. Second, we embed the same mechanism in a seven-sector model calibrated to China's 2020 input-output table. Real estate occupies a distinctive role in the multi-sector setting: it produces a durable asset held by households (rather than a flow good consumed by other firms), and households' housing demand responds endogenously to credit conditions in the real-estate sector through a time-varying preference weight χ_t calibrated to match the observed housing-price decline. Importantly, neither modifies the trade-credit transmission mechanism itself: the channel that propagates the shock upstream operates as in the two-sector model.

Calibrating the multi-sector model to Chinese data, we find that the trade-credit channel is quantitatively first-order. The baseline impulse response to a 25.9% tightening of real estate's borrowing limit—calibrated to match the observed contraction in developer bank credit between 2019 and 2022—generates a GDP trough of -4.7% and a cumulative 20-quarter loss of -31.9 percentage-point-quarters; holding the prepayment fraction at its steady-state value reduces this loss to -19.1, attributing about 40% of the aggregate response to the trade-credit channel (a $1.7\times$ amplification). The decisive evidence on the channel's quantitative importance comes from comparing the model's output to the magnitude implied by our empirical estimates. Aggregating the firm-level sales coefficient to the economy through a transparent identity—the Empirical Implied Loss Ratio (EILR)—we find that the model with endogenous trade credit reproduces the bulk of the empirically implied GDP loss, an out-of-sample success given that no aggregate moment enters its calibration. Complementary sectoral exercises show that the trade-credit channel dominates the upstream spillover. Policy counterfactuals further reveal that the aggregate response is shaped by the structure of the trade-credit chain rather than by the timing or intensity of the underlying credit shock.

This paper makes four contributions. First, we provide empirical evidence that trade credit

transmits sector-specific financial shocks through production networks. A central challenge in this literature is that comprehensive firm-to-firm trade credit data are rarely available; we address this by combining firm-level financial data with input-output network structure in a causal framework, showing that sector-specific credit shocks propagate broadly through the economy via interfirm payment relationships and generate significant real effects on cash holdings, sales, and asset values of exposed firms. Our research design further provides corroborating evidence from both sides of the credit relationshipdemand-side payment delays among directly affected firms and supply-side receivable accumulation among their upstream suppliersoffering unusually direct validation of the transmission mechanism. Among concurrent empirical assessments of the Three Red Lines policy specifically, the closest to ours is Chen et al. (2024), who use a stock-correlation-based exposure measure and document spillovers in stock prices, bond spreads, and corporate investment; our production-network-based exposure measure captures commercial rather than financial connections to real estate, and our two-sided demand–supply evidence is unavailable from a single-sided exposure design. Ma and Xie (2025) also studies this policy but focuses on developers' own financial risk rather than on transmission to non-real-estate firms.

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 as trade credit, and this fraction is pinned down by the shadow value of its binding borrowing constraint. Trade credit here is therefore not a separate credit market layered onto the network but an object generated by financial frictions on the network's own transaction links—which is exactly why it carries *financial*, rather than productivity, shocks upstream. This view complements two influential strands: Altinoglu (2021), who first brought endogenous trade credit into a production network, modeling it as supplier-provided credit collateralized by the downstream firm's revenue; and a cost-of-credit line (Reischer, 2020; Luo, 2020; Cun et al., 2022) that derives trade-credit terms from a choice between bank and supplier credit. Our buyer-side, binding-constraint lens nests these as facets of one object—the financing of input purchases—and supplies the channel that a regulatory *quantity* shock, one that tightens the borrowing limit rather than moving interest rates or collateral values, actually activates. It is also where our empirical heterogeneity points: the spillover concentrates in financially constrained firms.

Third, the framework is portable. Because it grounds trade credit in a binding constraint and

solves the network transparently, it applies to any setting in which a sector-specific financial shock meets a production network—regulatory deleveraging, bank-credit disruptions (Huremović et al., 2026), or distress at a network-central firm—offering a tractable tool beyond the present application.

Fourth, we provide a structural quantitative assessment of China's Three Red Lines and, with it, a structural–empirical bridge that is unusual in this literature: the model with endogenous trade credit reproduces the bulk of the empirically implied GDP loss, whereas without the channel it captures only about half—direct evidence that the endogenous adjustment of trade credit is quantitatively first-order. The broader lesson concerns the design of *sector-specific financial regulation*: a leverage cap aimed at a single sector propagates economy-wide through the trade-credit links of the production network, so such interventions must be evaluated for their network spillovers, not merely their effect on the targeted firms. For a sector as deeply embedded as Chinese real estate, most of the welfare cost runs through these spillovers rather than through the direct hit on developers.

The paper proceeds as follows. Section 2 provides institutional background on China's Three Red Lines policy and the real estate sector's position in the production network. Section 3 describes our data and presents the empirical analysis, documenting economy-wide spillover effects and the demand-side and supply-side mechanisms of trade credit transmission. Section 4 develops the theoretical framework—first a two-sector model illustrating the core mechanism, then a seven-sector multi-sector model that incorporates real estate as a final-use durable sector and an endogenous housing demand wedge—and reports calibration and the baseline impulse response. Section 5 presents the quantitative analyses: matching the model to empirical magnitudes, decomposing the trade-credit channel across sectors, and evaluating policy counterfactuals. Section 6 concludes with policy implications and directions for future research.

2 Institutional Background

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. Land-use-rights sales 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, Survey and Statistics Department, 2020), far exceeding the corresponding figure in most developed economies. This concentration of household wealth in real estate means that 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 development firms:

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 annual interest-bearing debt growth: *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 other industries' regulatory environment largely unchanged, allowing clean identification of spillover effects to non-real-estate firms.

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, real estate development investment 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

Understanding why the Three Red Lines policy generated upstream transmission requires some background on how Chinese real estate developers financed their operations. 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] 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—which constituted a substantial non-bank financing channel in the Chinese real estate sector.

In practice, developers routinely delayed payments to construction contractors, building-material

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

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 magnitude illustrates the extent to which large developers had come to rely on suppliers, contractors, and other counterparties as a source of 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 supply-chain financing channel 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 direct evidence of the transmission. 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 bil-

[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.

[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.

lion 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 Three Red Lines policy constrained developers' access to external financing, this pre-existing reliance on supplier financing became a channel through which developer liquidity stress was transmitted to upstream firms. 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.

3 Empirical Strategy and Results

3.1 Data

Our empirical analysis relies on three main data sources: firm-level financial data from listed companies, disclosed supply chain relationships, and input-output tables.

Firm-level financial data. We obtain quarterly financial statement data for Chinese listed firms from the China Stock Market & Accounting Research (CSMAR) database, covering the period from 2017Q1 to 2025Q3. Our primary variables include trade credit measures—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

[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.

(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.

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 tables. We obtain the 2020 input-output table from China's National Bureau of Statistics. This table provides industry-level data on intermediate goods flows 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 treatment intensity measures based on each industry's exposure to real estate through the production network.

3.2 Main Results: Economy-Wide Spillover Effects

We begin by establishing our central empirical finding: the Three Red Lines policy generated substantial economic spillovers that reached well beyond the targeted real estate sector.

3.2.1 Treatment Intensity

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

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.

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

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

The first component is the firm's average ratio of accounts receivable to total assets during the pre-policy period, capturing the extent to which the firm extends trade credit to its customers. Firms with higher AR ratios extend more credit to their customers and are therefore more vulnerable 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 $\Psi = (\mathbb{I} - \Omega)^{-1}$ corre-

sponding 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 trade credit disruptions. 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.

3.2.2 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 variation in firm-level trade credit exposure. Standard errors are clustered at the industry level.

3.2.3 Results

Table 2 reports our main estimates. The results indicate broad-based economic deterioration among firms more exposed to real estate through trade credit channels. A one-percentage-point increase in exposure is associated with a decline of 236 million RMB in cash holdings (column 1), 204 million RMB in quarterly sales revenue (column 2), 225 million RMB in interest-bearing debt (column 3), and 362 million RMB in equity (column 4). All 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 liq-

uid resources, consistent with delayed receipt of payments from real estate customers. The sales decline suggests that the financial shock translated into real activity contraction—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 capacity and equity implies a deterioration in firms' overall financial position, not merely a temporary cash flow disruption.

3.2.4 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.1 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 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 suffered liquidity losses, revenue declines, and balance sheet deterioration in proportion to their trade credit exposure to real estate. A natural question follows: through what mechanism did this transmission occur? We turn to this ques-

Table 2: Economy-Wide Spillover Effects

	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)
_cons	27.36***	32.29***	35.76***	71.39***
	(0.24)	(0.36)	(0.31)	(0.38)
N	142,024	139,429	142,024	142,024
Firm FE	Yes	Yes	Yes	Yes
Industry × Time FE	Yes	Yes	Yes	Yes
Adj. R^2	0.899	0.951	0.897	0.976

Standard errors clustered at industry level in parentheses.

Intensity = (Pre-policy AR/Assets) × (Leontief exposure to RE) × 100.

Estimates are robust to controlling for firm size (log total assets).

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

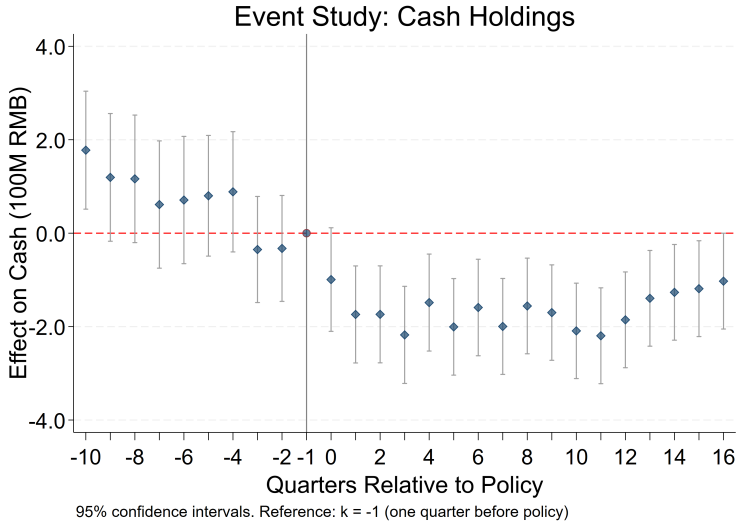


Figure 1: Event Study: Cash Holdings

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.

tion next, examining the channel from both the demand side (real estate firms' behavior) and the supply side (upstream suppliers' experience).

3.3 Demand-Side Mechanism: Real Estate Firms' Payment Behavior

For the trade credit channel to operate, a necessary condition is that the policy shock induced real estate firms to delay payments to their suppliers. We provide direct evidence on this "sending end" of the transmission.

3.3.1 Cross-Industry Evidence

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

$$Y_{it} = \beta \cdot (\text{RE}_i \times \text{Post}_t) + \theta_i + \lambda_t + \varepsilon_{it} \tag{3}$$

where RE_i indicates real estate development firms and Post_t equals one after 2020Q3. The specification includes firm and time fixed effects, with standard errors clustered at the industry level.

Table 3 reports the results. The estimates reveal a dramatic reallocation of real estate firms' financing sources following the policy. Accounts payable surged by 4.58 billion RMB, and the AP-to-sales ratio rose by 0.64 (column 2), indicating that real estate firms substantially increased the credit they obtained from suppliers by delaying payments. At the same time, cash holdings fell by 1.31 billion RMB and short-term bank borrowing contracted by 823 million RMB. The picture is clear: as bank credit dried up and internal liquidity depleted, 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 dynamic patterns underlying these average effects 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.

3.3.2 Within-Sector Variation

The cross-industry comparison, while informative, relies on comparing real estate firms to a heterogeneous set of other industries. To sharpen identification, we exploit the policy's tiered structure and compare firms within the real estate sector based on whether they breached any of the three regulatory thresholds (Red, Orange, or Yellow tier) versus those that met all thresholds (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.

Table 4 shows that the response was concentrated among firms most directly bound by the policy. Among real estate firms, those that breached regulatory thresholds increased accounts payable by 8.71 billion RMB relative to compliant firms, consistent with the more severe credit constraints they faced under the policy's tiered structure. Short-term borrowing fell by 1.01 billion RMB. 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 demand-side evidence establishes the first link in the transmission chain: the Three Red Lines policy constrained real estate firms' access to formal credit, forcing them to lean heavily on supplier financing. The accounts payable of real estate firms—which are the accounts receivable of their suppliers—surged, transmitting the liquidity shock upstream through the production network.

3.4 Supply-Side Mechanism: Evidence from Upstream Suppliers

The demand-side analysis shows that real estate firms delayed payments, but this does not by itself prove that suppliers were adversely affected—firms might have had sufficient buffers to absorb the delayed receipts. We now examine the "receiving end" of the trade credit channel, asking whether upstream suppliers experienced financial deterioration as a result.

Table 3: Demand Side: RE Firms' Payment Behavior

	(1) AP (100M)	(2) AP/Sales (Ratio)	(3) Cash (100M)	(4) ST Loan (100M)
RE × Post	45.77***	0.64***	-13.10***	-8.23***
	(5.45)	(0.09)	(1.79)	(0.87)
N	142177	139571	142177	142177
Firm FE	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
Adj. R^2	0.847	0.152	0.894	0.826

Standard errors clustered at industry level in parentheses.

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

Table 4: Within-RE Heterogeneity: Breached vs Green Firms

	Dependent Variable (100M RMB)		
	(1) AP	(2) ST Loan	(3) LT Loan
Breached × Post	87.12**	-10.05**	26.08
	(34.07)	(4.47)	(21.45)
N	3856	3856	3856
Firm FE	Yes	Yes	Yes
Time FE	Yes	Yes	Yes
Adj. R^2	0.825	0.740	0.912

Standard errors clustered at firm level in parentheses.

Sample: RE firms (K70). Treatment: Breached ≥1 threshold vs. Green.

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

3.4.1 Evidence from Disclosed Supplier 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 had 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:

$$Y_{it} = \beta \cdot (\text{Treat}_i \times \text{Post}_t) + \theta_i + \gamma_t + \varepsilon_{it} \tag{5}$$

To ensure comparability, we restrict the sample to a balanced panel of firms observed both before and after the policy.

Table 5 presents the results. Firms with real estate customers saw their accounts receivable rise by 9.7 percentage points of sales after the policy (column 1), consistent with the payment delays by real estate firms translated into higher receivables at their suppliers. The real consequences are substantial: log sales fell by 0.21 (column 2), log assets by 0.11 (column 3), and total debt by 2.68 billion RMB (column 4). These suppliers not only absorbed the credit shock passively through rising receivables but also experienced operational contraction—lower sales and shrinking balance sheets—consistent with the liquidity squeeze propagating into real activity. 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.

3.4.2 Evidence from Input-Output Linkages

The disclosed-supplier analysis provides compelling micro-level evidence but covers only a small sample of firms with voluntarily reported customer relationships. To assess whether the transmission pattern holds more broadly, we exploit industry-level input-output linkages across the full sample of non-real-estate listed firms:

$$Y_{it} = \beta \cdot (\Psi_{kh} \times \text{Post}_t) + \theta_i + \lambda_t + \varepsilon_{it} \tag{6}$$

Table 5: Supply Side: Disclosed RE Suppliers

	(1) AR/Sales (Ratio)	(2) ln(Sales)	(3) ln(Assets)	(4) Debt (100M)
Treat × Post	0.097*	-0.208**	-0.106*	-26.836**
	(0.055)	(0.103)	(0.057)	(13.249)
N	2547	2565	2566	1488
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Adj. R^2	0.498	0.937	0.962	0.955

Cols. (1)–(3) clustered at industry level; col. (4) at firm level.
Treatment: firms with disclosed RE customers (2018–2020) vs. without.
Annual data, balanced panel. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

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.

Table 6 reports the results. Among trade credit variables, bills receivable increase significantly (column 2), though accounts receivable alone is imprecisely estimated (column 1)—a pattern likely reflecting the prevalence of commercial acceptance bills as the dominant instrument for deferred payment in China's real estate supply chain. On the real outcomes side, industries more linked to real estate through the production network experienced significant declines in cash holdings (478 million RMB per unit of exposure), sales (498 million RMB), and total assets (3.00 billion RMB). These effects corroborate the micro-level findings from disclosed relationships, indicating that trade credit transmission operated at scale across the broader economy.

The supply-side analysis validates the mechanism from both ends of the trade credit relationship. The micro-level evidence from disclosed customer-supplier links shows that firms with real estate customers experienced a sharp rise in their accounts receivable relative to sales, alongside declines in sales, assets, and debt. The macro-level evidence from input-output linkages shows that industries more connected to real estate through the production network experienced significant increases in bills receivable and broad declines in real outcomes. The two approaches—one using narrow but precise firm-pair links, the other using comprehensive but coarser industry linkages—deliver mutually corroborating evidence of supplier exposure.

Table 6: Supply Side: IO-Based Exposure

	Trade Credit Variables			Real Outcomes (100M RMB)		
	(1) AR (100M)	(2) BR (100M)	(3) TCS (Ratio)	(4) Cash	(5) Sales	(6) Assets
$\Psi_ih \times 100 \times$ Post	-1.84	0.91***	0.04	-4.78***	-4.98*	-30.04**
	(1.51)	(0.30)	(0.25)	(1.50)	(2.97)	(11.98)
N	138321	138321	135739	138321	135749	138321
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R^2	0.860	0.473	0.161	0.894	0.947	0.949

Standard errors clustered at industry level in parentheses. RE sector excluded.

TCS = (AR + BR + Prepayments)/Sales. Ψ_ih: Leontief inverse from industry $toRE$.

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

Mechanism dynamics, heterogeneity, and real-economy outcomes (Appendix B). 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.2). 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.3). Third, the absorption is highly uneven across firms: private and smaller firmsthose 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.4). Finally, a continuous dose-response specification confirms a monotone gradient across the intensity distribution (Appendix A.1). Together these auxiliary results already point toward an interpretation of the channel as fundamentally financialoperating through firms' working-capital positions, and concentrated where alternative financing margins are unavailable. The next section 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.

3.5 Mechanism III Network Position and Financial Constraint

Sections 3.3–3.4 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 2020-01-01 ownership, with private firms treated as more constrained than SOEsa distinction we verify delivers consistent results when proxied by firm size (Appendix A.4). Table 7 estimates the channel by network order, ownership, and 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 estatedominated

in China by finance, business services, and utilitiesdo not transmit the shock upstream, plausibly because their alternative liquidity margins absorb the demand contraction without forcing them to squeeze their own suppliers. Within the chain of higher-order suppliers that does carry the channel, financial constraint amplifies absorption: 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 ownthe effect emerges multiplicatively, only where the chain reaches *and* the receiving firm cannot substitute around the involuntary lending.

This signatureindirect chain propagation interacting with firm-level constraintadds structural detail to the basic finding of Sections 3.3–3.4. 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, and in particular to the contrast with one-dimensional alternatives, in the empirical summary that follows.

4 Model

The empirical analysis in Section 3 establishes three facts about how China's Three Red Lines policy transmitted through the production network. First, firms more exposed to real estate through trade-credit linkages experienced broad-based financial deterioration, with declines in cash, sales, and assets that persist for over three years. Second, constrained real estate firms responded by delaying payments to suppliers, with accounts payable rising sharply in the post-policy period. Third, disclosed suppliers of real estate developers absorbed this pressure through rising accounts receivable and contracted sales. To formalize this transmission mechanism and make it quantitatively useful, we develop a structural model in two steps.

Section 4.1 presents a two-sector model that makes the core mechanism transparent: a tightening of one sector's bank-credit constraint induces the constrained firm to reduce prepayment to its supplier, which transmits the financing pressure upstream through endogenous trade credit. Section 4.2 embeds the same mechanism in a seven-sector production network calibrated to China's 2020 input-output table. Real estate plays a distinctive role in our setting: it produces a durable asset held by households (rather than a flow good consumed contemporaneously by other sectors), and household demand for housing varies endogenously with credit conditions in the real-estate

sector through a time-varying preference wedge. Both features capture the role of new housing as a final-use asset rather than a productive intermediate; neither modifies the trade-credit transmission mechanism, which operates identically in both versions of the model. Section 4.3 describes the calibration and Section 4.4 reports the baseline impulse response. Section 5 reports the quantitative analyses—matching empirical magnitudes, the sectoral decomposition of the trade-credit channel, and a set of policy counterfactuals.

4.1 A Two-Sector Model with Endogenous Trade Credit

This subsection isolates the trade-credit transmission mechanism in the simplest possible setting. We strip the model to two production sectors and a single period, taking factor prices, the intermediate-good price, the downstream output price, and interest rates as given. This partial-equilibrium device deliberately shuts down the general-equilibrium price and demand feedbacks that the multi-sector model in Section 4.2 captures; its only purpose is to display, analytically, how a financial shock to one sector reaches its supplier through endogenous payment terms. The quantitative magnitude of the channel—including its amplification of the aggregate response—is the task of the calibrated multi-sector model.

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 we

document empirically, 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 which the transmission concentrates among financially constrained firms rather than tracking interest-rate differentials.

From two sectors to the network. The two-sector model is deliberately confined to a single transmission link and a fixed-price environment, which makes the mechanism transparent but understates its aggregate footprint. The multi-sector model of Section 4.2 preserves the endogenous-prepayment condition (8) sector by sector and embeds it in the 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 documented in Proposition 1 into quantitatively important amplification of the aggregate shock.

4.2 Multi-Sector Model for the Chinese Economy

The two-sector model 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, household's 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 $L_t^{\eta} =$

$W_t/(P_tC_t)$ and the housing Euler equation

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

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 $S^C_{i,t} = \prod_{j \neq i} P_{j,t}^{\zeta_{ij}}$ and the within-sector intermediate demand satisfies $P_{j,t} X^C_{ij,t} = \zeta_{ij} S^C_{i,t} X_{i,t}$.

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 $G(\phi_t) S^C_t X^C_t$ with $G(\phi_t) = \frac{1}{2}(1 - \phi_t)^2$. The firm's financing constraint requires that net cash outflows be financed by bank credit subject to the limit θ_i:

$$\underbrace{\phi_{i,t} S^C_{i,t} X^C_{i,t}}_{\text{prepayment to suppliers}} + \underbrace{W_t L_{i,t}}_{\text{wages}} - \underbrace{\sum_{j \neq i} \phi_{j,t} \zeta_{ji} S^C_{j,t} X^C_{j,t}}_{\text{trade credit received from customers}} - \underbrace{\phi^H_t \zeta_{hi} S^H_t X^H_t}_{\text{trade credit received from RE}} \leq \theta_i P_{i,t} Y_{i,t}. \tag{14}$$

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

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

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

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

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 model: 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 the commodity sectors,

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

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 transmission channel that is the central focus of the paper.

4.2.4 Endogenous Housing Demand Wedge χ_t

The household's housing Euler equation (13) implies that, under a constant preference weight, a tightening of θ_t^H generates a counterfactual response in Q_t: IH_t falls, H_t falls below trend, the term χ_t / H_t rises, and Q_t rises with it. This implication contradicts the data: China's 70-city new-home price index declined by approximately 2.3% over the eight quarters following 2020Q3.

To bring the model into line with the observed price response, we let χ_t evolve as an AR(1) driven by the credit shock:

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

where $\kappa_\chi > 0$ controls the long-run elasticity of χ_t with respect to the real-estate borrowing-constraint shock and $\rho_\chi \in [0, 1)$ controls persistence. We treat (21) as a reduced-form summary of three real-world forces. The first is delivery risk on pre-sold housing: financially distressed developers fail to complete pre-sold projects, an issue documented extensively in 2021–2024 China (the "unfinished housing" phenomenon). The second is secondary-market liquidity disruption: tighter developer credit reduces the depth of the resale market for housing held by households. The third is a revision of long-run housing valuation expectations, as the policy signals a regime change in the government's stance toward real estate. None of these channels is explicitly micro-founded in the model; χ_t stands in for their net effect on household demand. The distinction relative to standard preference shocks in the housing macro literature (Iacoviello, 2005; Liu et al., 2013) is that our χ_t is linked endogenously to the underlying credit shock, rather than introduced as an independent disturbance.

A key conceptual point: 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. The trade-credit transmission mechanism that is the central focus of our paper 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.

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.

4.3 Calibration

We calibrate the model at quarterly frequency to capture the Chinese economy at the eve of the Three Red Lines policy (2018Q1–2020Q2). The calibration draws on three data sources: China's 2020 national input-output table at the 153-sector level (aggregated to seven sectors); CSMAR firm-level financial statements; and standard macro-finance parameters from the literature. We organize the discussion into four blocks — preferences and depreciation, the production-network structure, 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 annual real interest rate of approximately 4%. The inverse Frisch elasticity is $\eta = 0.5$, in the middle of standard ranges used in housing-cycle DSGEs. The quarterly housing depreciation rate is $\delta_h = 0.01$, consistent with an annual depreciation of approximately 4% used for residential housing in Iacoviello and Neri (2010). The intertemporal elasticity of substitution is unity ($\sigma = 1$, log utility over C). The steady-state housing preference weight $\bar{\chi} = 0.40$ is chosen so that the SS housing-wealth-to-consumption

ratio $QH/(PC)$ is in line with Chinese household balance-sheet aggregates documented in Liu et al. (2013).

4.3.2 Production-Network Structure

We aggregate the 2020 IO 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 IO 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 intermediate-input intensity.

Consumption shares. Final consumption shares γ_i are computed from the IO 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 IO 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 IO 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 sectoral financial parameter. 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 SS constraint multipliers μ_i are strictly positive but moderate, which ensures well-behaved IRF dynamics.

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 SS trade credit ($\bar{\phi}_i = 1 - R_i - \bar{\mu}_i$) but has limited influence on the IRF dynamics, 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 observed cumulative contraction in real estate developers' bank-loan-to-revenue ratio 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 housing-preference-shock literature (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 IO-to-7-sector aggregation map are deferred to Appendix D.

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 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

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.

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; the cumulative 16-quarter loss is -30.9 percentage points-quarters, which is the central quantity the counterfactual decomposition in Section 5 uses to isolate the trade-credit channel.

Figure 3 reports the IRF 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 FOC (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 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.

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.

5 Quantitative Analysis

5.1 Matching Model Predictions to Empirical Magnitudes

The reduced-form estimates in Section 3 identify a firm-level relationship between trade credit exposure and post-policy outcomes, but they speak in different units from the structural model: 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. To quantify whether the model's predicted GDP loss is consistent with the magnitude implied by the empirical estimates, we construct an Empirical Implied Loss Ratio (EILR) and compare it directly to the model's IRF under the two counterfactual scenarios introduced above.

Construction of the EILR. Let β_S denote the "Sales" coefficient in column (2) 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

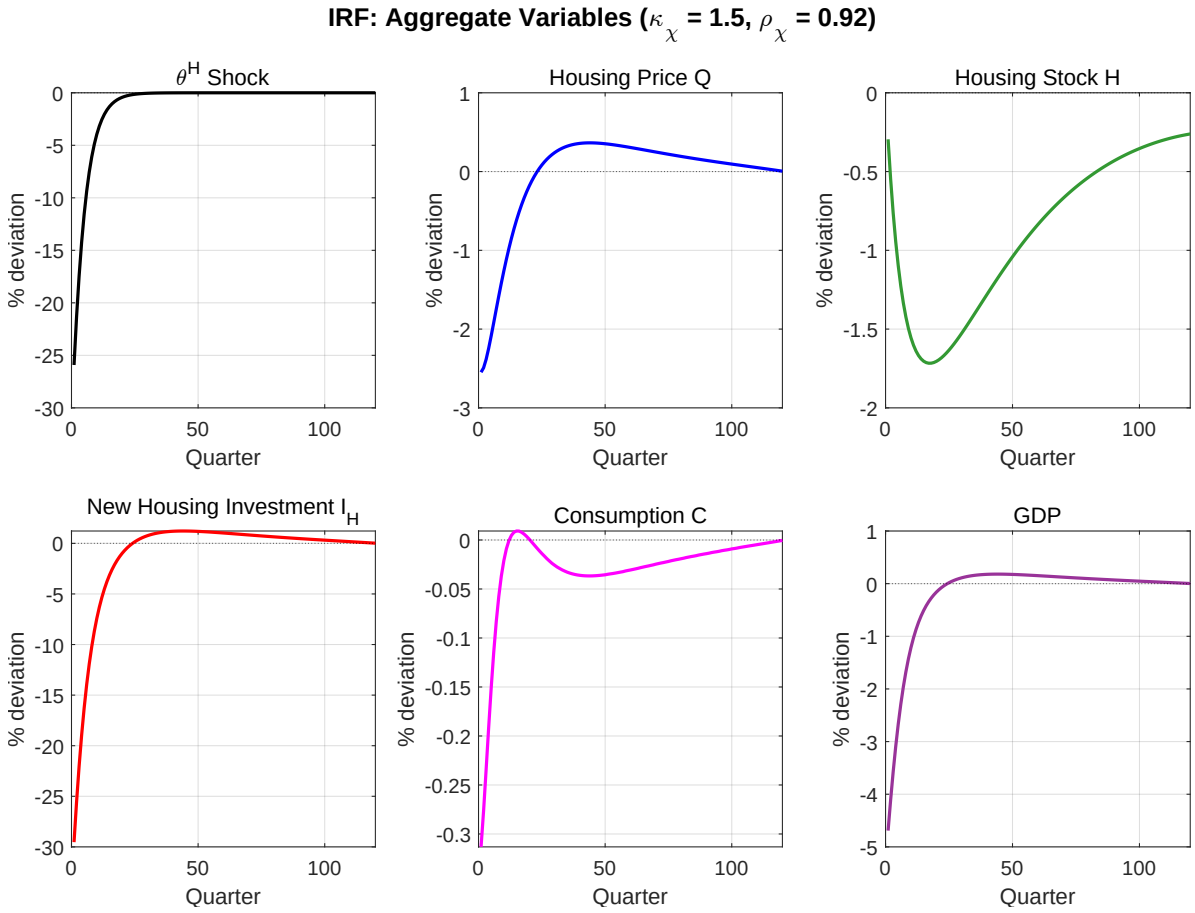


Figure 2

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 $(\kappa_\chi, \rho_\chi) = (1.5, 0.92)$.

effect across firms and scaling to economy-wide GDP yields a per-quarter implied GDP loss

$$\text{EILR}_q \;=\; \frac{\beta_S \cdot \bar{I} \cdot \lambda}{\bar{S}} \times 100, \qquad \text{EILR}_{\text{cum}} \;=\; Q_{\text{post}} \cdot \text{EILR}_q, \tag{22}$$

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 (22) a function of only four sample moments and one regression coefficient.

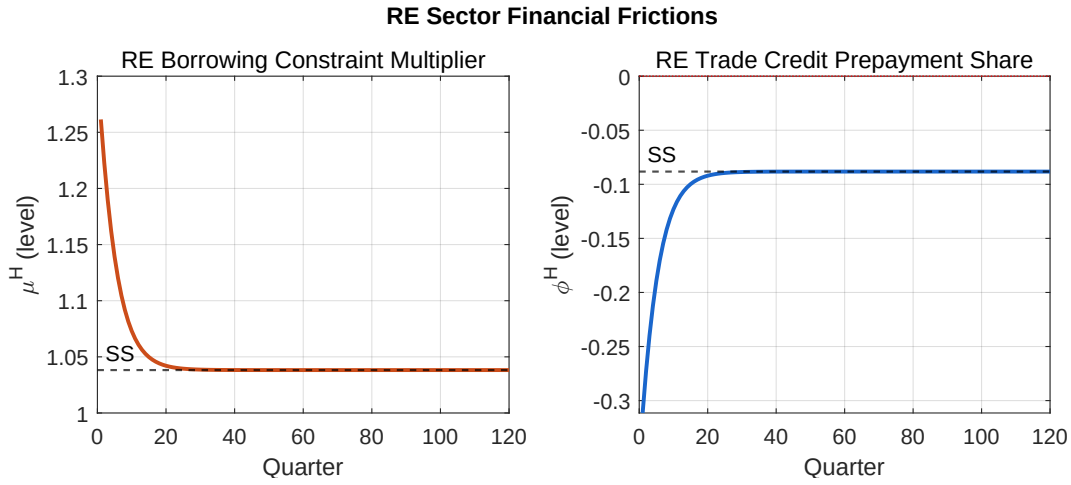


Figure 3

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.

Inputs. We obtain $\beta_S = -2.04$ (column (2) 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 (22) 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 compares the EILR with the model's 20-quarter cumulative GDP deviation under the two scenarios from the counterfactual decomposition. Two findings stand out.

First, the trade-credit channel is quantitatively first-order, and this is the result we emphasize. The cumulative GDP loss in Scenario A (endogenous ϕ) is -31.9 pp·q; holding the prepayment share ϕ at its steady-state value and severing it from the borrowing constraint (Scenario B) reduces the loss to -19.1 pp·q. The endogenous adjustment of trade credit thus accounts for about 40% of the aggregate response—a 1.7× 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

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×.

empirical estimate.

Second, the model reproduces the bulk of the empirical magnitude. Scenario A's −31.9 pp·q is about four-fifths of the empirically implied loss of −40.4 pp·q, whereas Scenario B—the model a researcher would obtain by treating trade credit as exogenous, as in Altinoglu (2021)—captures only about half. Because β_S enters nowhere in the calibration of θ_h, ν_i, or the production network, the agreement under Scenario A is an out-of-sample success: 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. This benchmark is a transparent order-of-magnitude bridge, 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 highly persistent—roughly flat over the post-policy window. The cumulative *match level* is therefore sensitive to the comparison window and to the model's calibrated shock persistence, even though the channel decomposition (the 40% share and 1.7× amplification) is not. We accordingly read the match level as an order-of-magnitude check and emphasize the window-invariant decomposition. (ii) The scaling by λ assumes listed-firm responses are representative of the economy, which may overstate the aggregate response if unlisted firms

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).

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). None of these caveats overturns the central message: within this transparent benchmark, the trade-credit channel raises the model's predicted loss by roughly 70% and brings it from about half of the empirically implied magnitude to the bulk of it.

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

The aggregate result of the previous subsection masks a striking sectoral heterogeneity: 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 counterfactual scenarios 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.

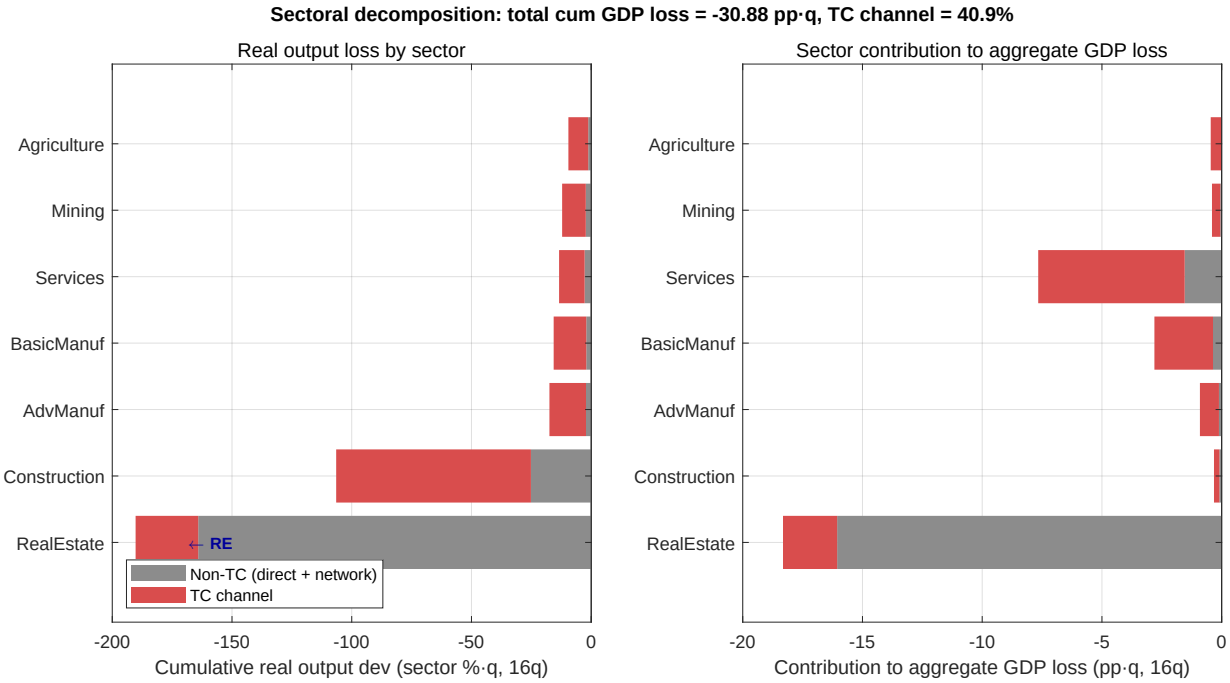


Figure 4

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.

The trade credit channel dominates spillover to upstream sectors. Across the six non-RE sectors, the trade credit channel accounts for between 76.4% and 87.5% of cumulative real output decline. 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_B^i 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). The Leontief production network by itself transmits very little of the regulatory shock past real estate's boundary; the upstream spillover that the empirical section identifies is structurally a trade credit phenomenon.

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 es-

tate 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 via 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 RE'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: services experience only a mild per-unit decline, but the sector's 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 trade-credit channel's aggregate footprint is dominated by services. Services alone contribute -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 flows through a sector that does not feature prominently in supply-chain narratives, simply because that sector is large enough that an 80% TC share applied to a -0.8% quarterly shock generates more aggregate impact than a 90% TC share applied to a -2.5% quarterly shock in a smaller sector.

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 cum$^i_B = -25.2$ sector %·q — roughly ten times larger in absolute value than the corresponding

figure for any other upstream sector — captures the Leontief amplification of RE's fall in input demand, before any trade credit spillover. 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 empirical evidence in Section 3 that AR/Assets-weighted Leontief exposure 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 seven sectors all 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 with different Leontief distance 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 on 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 policy 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: 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 **??** (Scenario B). Numbers reported to two decimals; values may shift slightly across runs because of fsolve convergence tolerance.

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

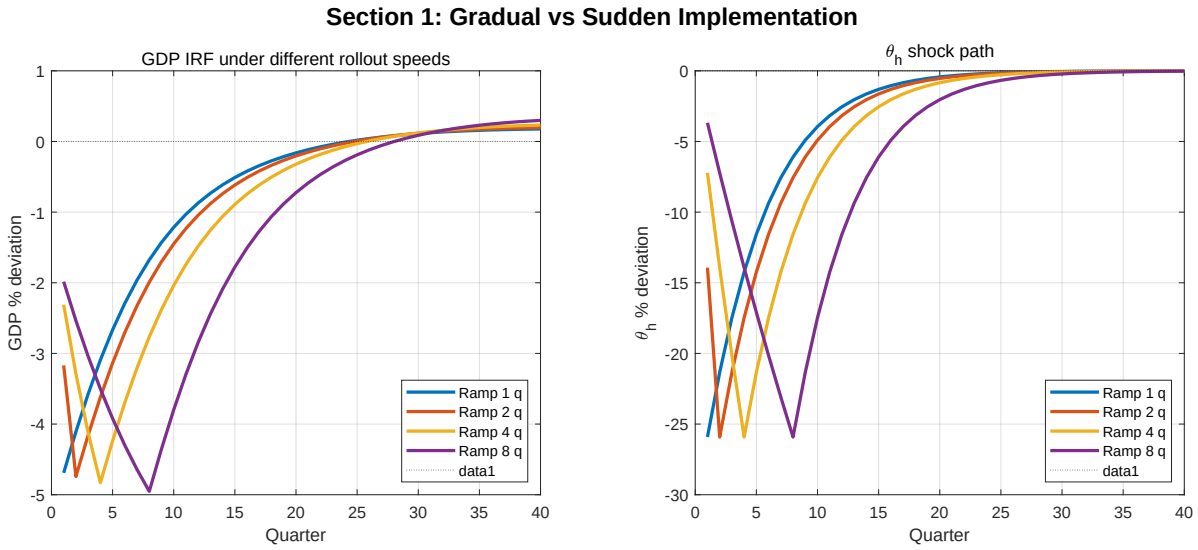


Figure 5

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%.

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 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.

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 di-

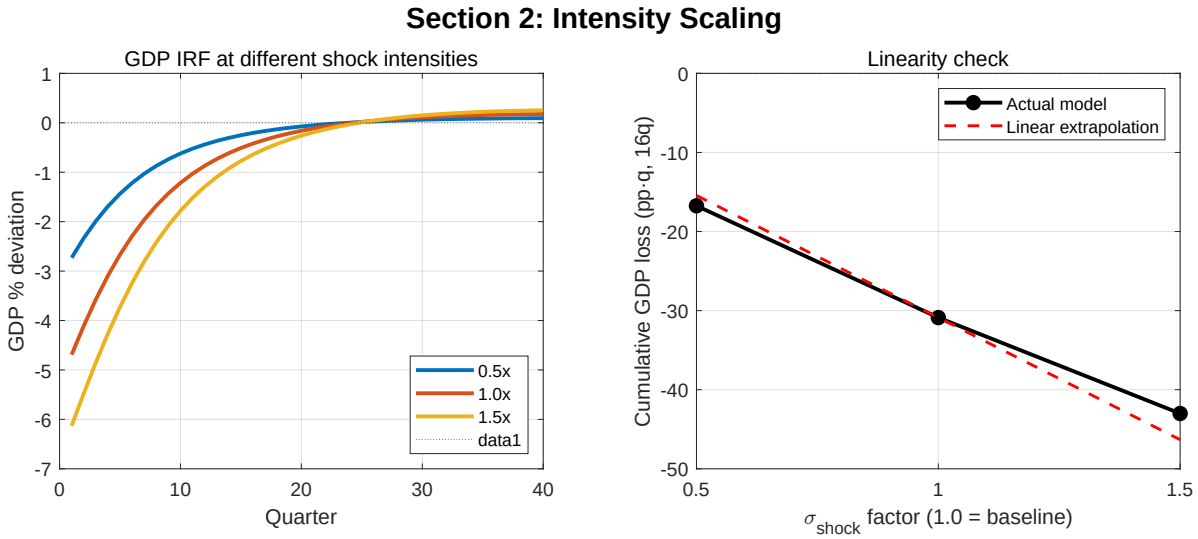


Figure 6

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).

minishing amplification at the margin. Over the range we consider, the trade-credit channel is approximately—though not exactly—linear in shock magnitude.

5.3.3 Trade-Credit Regulation

The third exercise asks whether a complementary regulation on 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 $\rho \in \{0, 0.25, 0.5, 0.75, 0.9, 1\}$ of the steady-state value $\bar{\phi}^H$.

A brief note on the steady-state level of ϕ^H is useful before turning to the results. 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 trade-credit extraction 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 intervention: it requires real estate to pay suppliers at least at par, ruling out the steady-state trade-credit extraction. 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 on 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 pres-

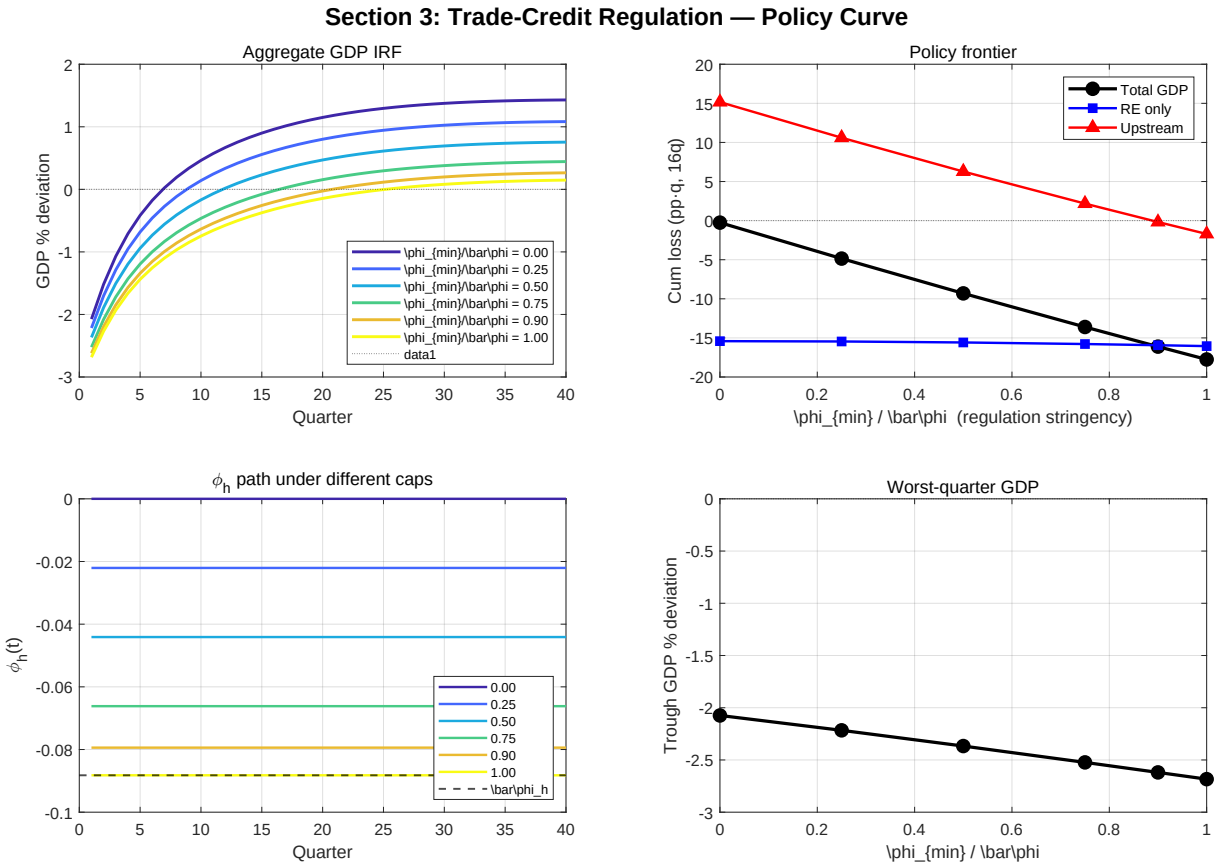


Figure 7

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.

sure 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 trade-credit substitution channel 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 accounts-receivable write-downs 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 anticipated as part of the policy's design horizon.

6 Conclusion

This paper studies how sector-specific financial shocks propagate through production networks via trade credit, providing both empirical evidence and a theoretical framework. Exploiting China's 2021 "Three Red Lines" policy as a quasi-natural experiment, we document that firms more exposed to real estate through trade credit channels experienced significant declines in cash holdings, sales, and assets, with effects persisting for over three years. Evidence from both sides of the trade credit relationship—rising accounts payable at constrained real estate developers and rising accounts receivable at their suppliers—identifies trade credit as the operative transmission channel for the spillover. To formalize this mechanism, we develop a tractable dynamic production network model in which trade credit terms are endogenously determined by firms' financing decisions. Calibrating the model to China's input-output structure, we find that endogenous trade credit adjustment accounts for approximately 41% of the policy's cumulative GDP impact and brings the model from 56% to 95% of the empirically implied loss—a direct structural-empirical validation of the channel's quantitative importance.

The central insight is that trade credit transforms what would otherwise be a localized regu-

latory intervention into an economy-wide disturbance. When credit-constrained firms delay payments to their suppliers, they effectively convert their own liquidity shortage into a credit supply contraction for their trading partners. In dense production networks, this mechanism propagates distress far beyond the targeted sector—through linkages that are invisible to standard input-output analysis but 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 for sectors as deeply embedded in production networks as Chinese real estate, the welfare 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 anticipated as part of any sector-specific intervention's design horizon, and direct-impact assessments that focus only on the targeted firms substantially underestimate the aggregate response. For researchers, our framework provides a portable tool for studying other settings where sector-specific financial shocks interact with production networks—including banking sector regulation, sector-targeted industrial policy, and credit market interventions in other emerging economies. The methodology of combining firm-level trade credit reliance with network-derived exposure measures, developed here to overcome data limitations on bilateral credit flows, is similarly applicable to other contexts where such bilateral data are unavailable.

We close with two caveats. First, the multi-sector application customizes the framework for the housing setting through two elements: 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 customizations rather than features of the underlying trade-credit framework. The core mechanism—endogenous trade credit responding to binding bank-borrowing constraints in a dynamic production network—does not depend on them and transfers directly to other settings where sector-specific financial regulation interacts with production networks. Second, our empirical analysis is restricted to listed firms, which may underrepresent the small and medium enterprises most likely to be affected by trade credit shocks;

richer firm-level data could reveal heterogeneous effects across the firm size distribution.

Several extensions of this work are worth pursuing. The framework can be applied to other episodes of sector-specific financial regulation—banking sector deleveraging, energy sector restrictions, or trade policy shocks—to assess whether trade credit amplification is similarly important elsewhere. The endogenization of trade credit could be enriched by introducing 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 are not captured by standard production network or financial accelerator frameworks. Building richer models of these networks—and gathering the bilateral firm-pair data needed to discipline them—is a promising direction for future research.

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.

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.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna, "Difference-in-Differences with a Continuous Treatment," *NBER Working Paper No. 32117*, 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.

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.

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.

Huremović, Kenan, Gabriel Jiménez, Enrique Moral-Benito, José-Luis Peydró, and Fernando Vega-Redondo, "Production and Financial Networks in Interplay," *American Economic Review*, May 2026, *116* (5), 1611–1647.

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.

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, Xin and Huobao Xie, "Real Estate Policy Regulation and Corporate Financial Risk: China's Three Red Lines Policy," *Pacific Economic Review*, 2025, *30* (1), 46–87.

People's Bank of China, Survey and Statistics Department, "2019 Survey on Assets and Liabilities of Urban Households in China," China Finance 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.1) verifies that the channel operates monotonically across the intensity distribution rather than only at extreme exposure. The direct trade-credit dynamics (A.2) document the payment-delay mechanism in the data. The real-economy outcomes (A.3) trace the channel beyond balance sheets to investment, inventory, and profitability. The financial-constraint heterogeneity (A.4) provides the full set of outcomes and constraint proxies that motivate the joint test in Section 3.5. The remaining sections report the full network decomposition across outcomes (A.5), compare the ecological treatment intensity to a direct dyadic measure constructed from Wind disclosure data (A.6), and explore aggregation to NBS industry value-added (**??**).

A.1 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,

$$Y_{it} \;=\; \sum_{k=2}^{5} \beta_k \cdot \big(\mathbf{1}\{Q_i = k\} \times \text{Post}_t\big) \;+\; \theta_i \;+\; \gamma_{kt} \;+\; \varepsilon_{it}. \tag{23}$$

A monotone decline of β_k in k is the testable dose-response prediction. Table 13 and Figure 8 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.

Table 13: 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. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

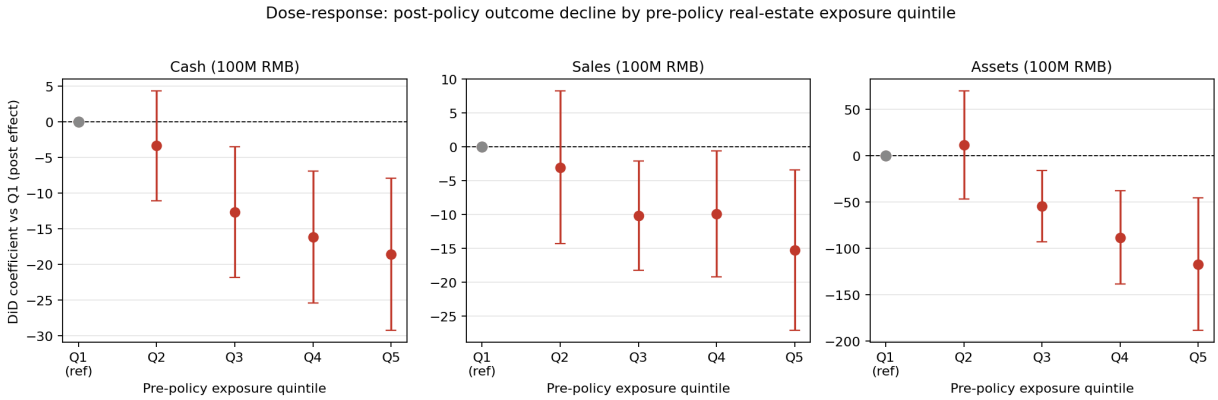


Figure 8: Dose-response by pre-policy real-estate exposure quintile

Notes: Each panel plots the quintile-specific post-policy effect (β_k in (23)) 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.

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 26). 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 partial recovery.

Table 14: 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. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

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 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

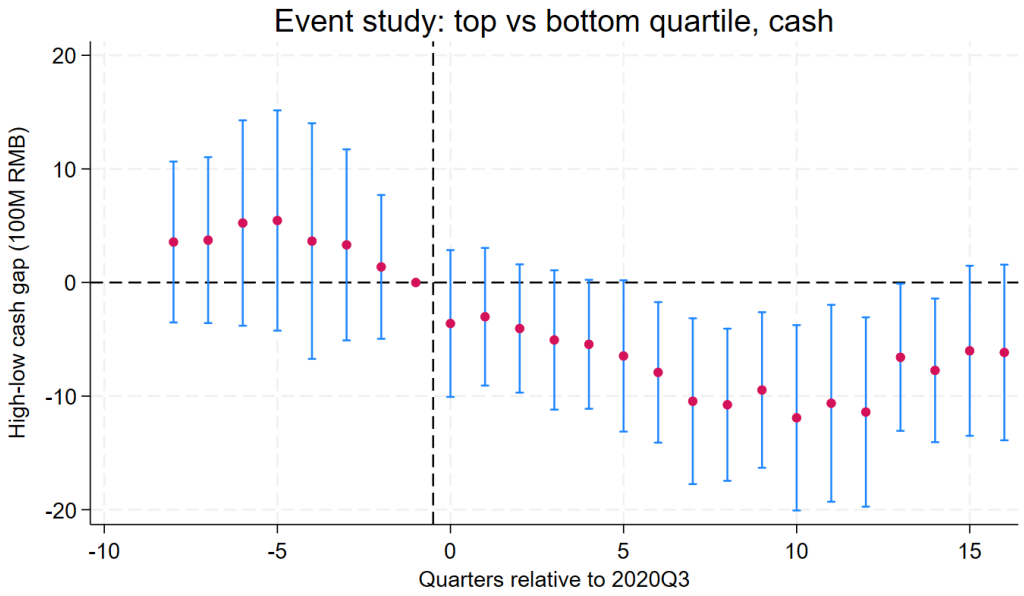


Figure 9: Event study: cash holdings, high vs. low exposure (median cutoff)

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.

causally.

A.2 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 15 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 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.

Table 15: 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.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

A.3 Real-Economy Outcomes

Beyond financial balance sheets, the channel propagates to real production decisions (Table 16). 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 16: 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.4 Heterogeneity by Financial Constraint

Section 3.5 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.

Table 17: 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.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

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

Table 18: 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.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

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.5 Network Decomposition: Full Results

Section 3.5 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 19 reports the network decomposition on cash with finer-grained components: direct (A), second-order (A^2), and the cumulative higher-order residual. Table 20 repeats the analysis on sales. Table 21 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-

Table 19: 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 20: 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 21: 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.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

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.

A.6 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 22 implements this dyadic approach using Wind's top-five AR-debtor disclosure (5,456 firms, 2018–2020 pre-policy snapshot); Table 23 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 indistin-

Table 22: 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 23: 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.

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

guishable 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 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 $\alpha \leq \frac{1}{2}$ (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 24 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.

Table 24: 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.

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 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

Table 25: 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).

after the within-RE entry is zeroed.

Table 25 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 $\zeta_{i,h}$).

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 (22) 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 (22). 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 $1.7\times$ 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.

Table 26: 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. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

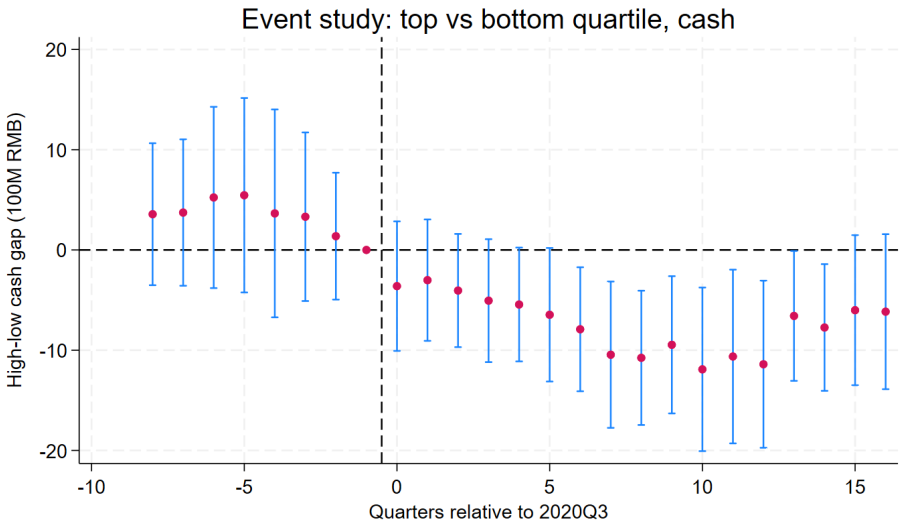


Figure 10: High-vs-Low Event Study: Cash

Notes: Event-study coefficients on (top-quartile × period dummies) for cash holdings, with the period one quarter before the policy ($k = -1$) as the reference and the policy quarter $k = 0$ marking 2020Q3. Sample restricted to firms in the top and bottom quartile of pre-policy intensity. Bars are 95% confidence intervals; standard errors clustered at the industry level.

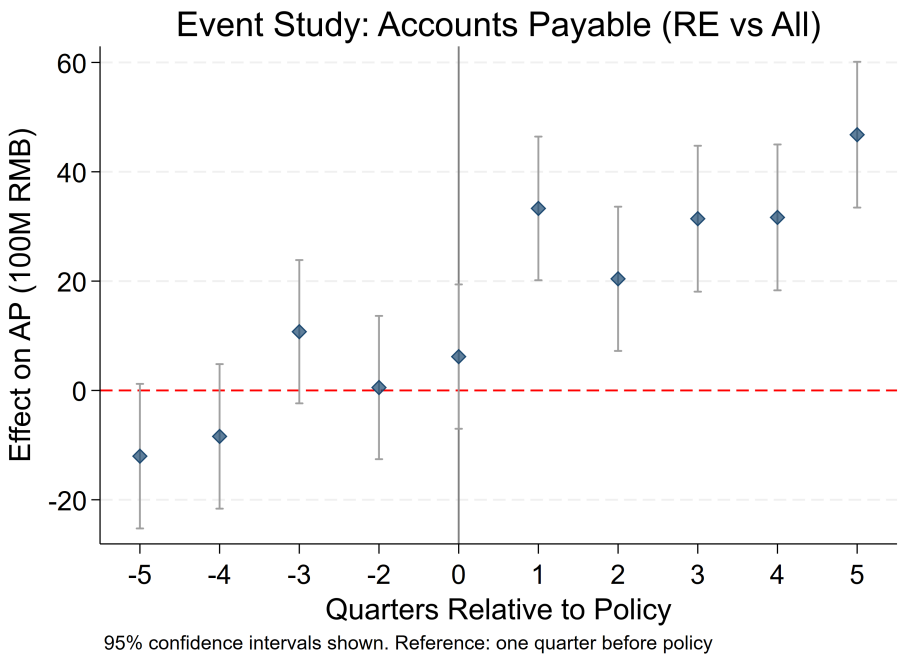


Figure 11: Event Study: Accounts Payable (RE vs. all)

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

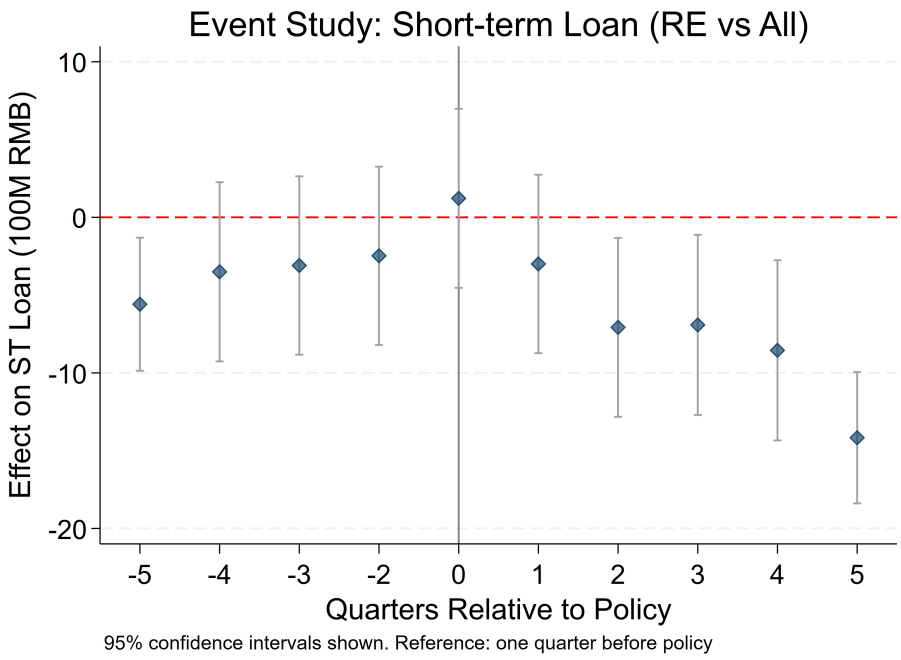


Figure 12: Event Study: Short-term Loan (RE vs. all)

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

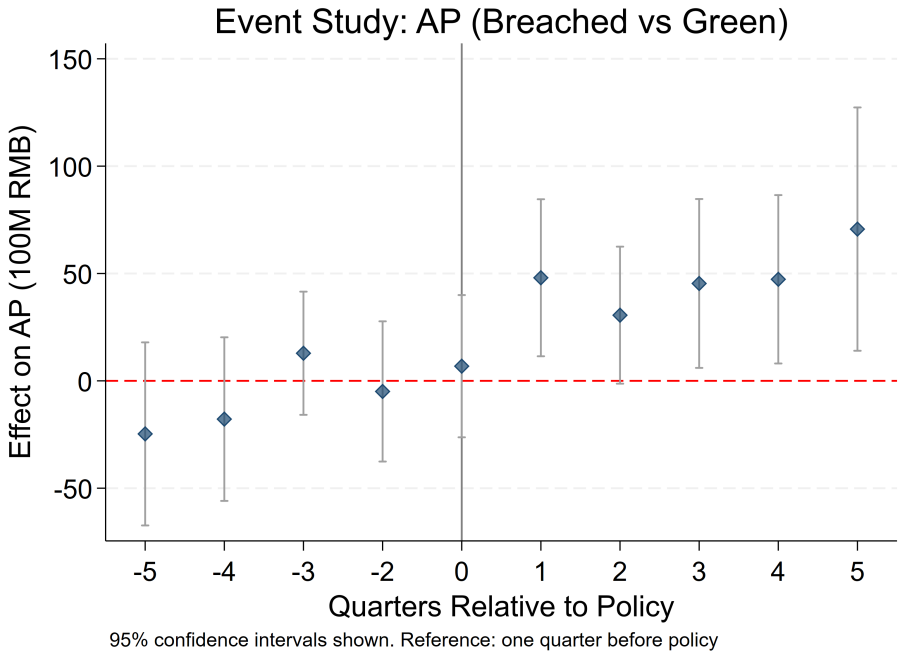


Figure 13: Event Study: AP (Breached vs. green within RE)

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.

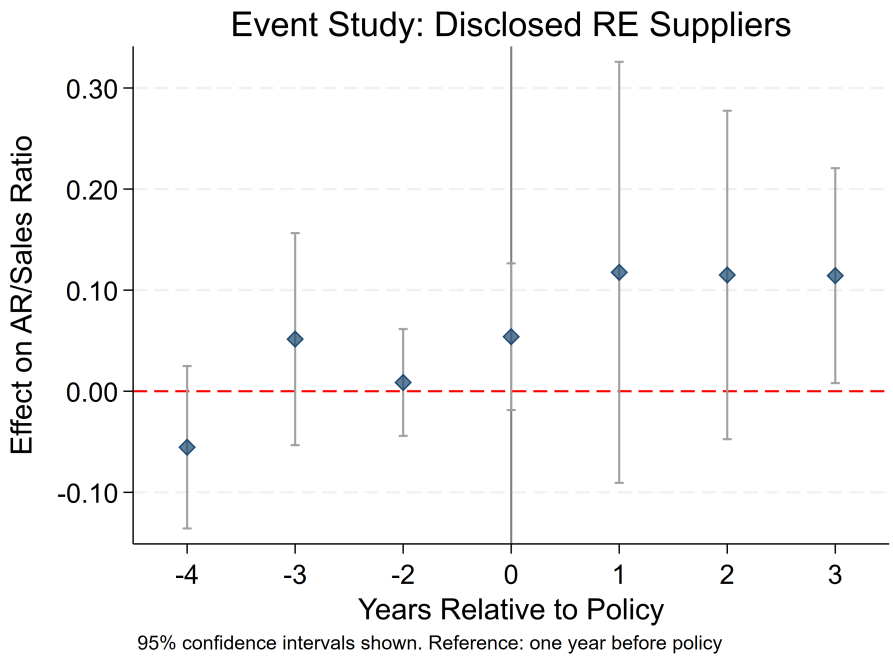


Figure 14: Event Study: AR/Sales (Disclosed RE suppliers)

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.