The Propagation of Environmental Risk Through Production Networks: Borrowing Cost Effects

Alessandro Dario Lavia *
University of Torino

Elisa Luciano
University of Torino & CCA

This Draft: February 22, 2026

[Click here for the most recent version]

Abstract

The cost of debt for firms more exposed to transition risk has been shown to be higher than the one of less exposed firms. To explain this difference, we develop a general equilibrium model featuring firms, banks and households. Firms are connected by a production network. Their input choices are rigid, in the sense that their decisions cannot be fully adapted to the actual risk realization. As a consequence of rigidity, default may occur. Interest rates on corporate debt are determined by banks, taking into account possible defaults. Without connections, brown firms would face higher costs because of transition exposure. In a network, the exposure depends also on each firm's connections. Even low-emission firms are penalized when they depend on brown suppliers, as upstream risks are transmitted through the supply chain. Debt costs depend on the overall exposure to transition risks. Empirically, we proxy firms' exposure to transition risk using sectoral CO_2 emissions and construct a network-based measure of total, embodied emissions that is fully consistent with the theoretical framework. Combining U.S. firm-level financial data with EPA emissions and input–output linkages, we show that lenders price both direct emissions and, crucially, network-adjusted carbon exposure, in line with the model's predictions. Evidence surrounding the Paris Agreement confirms the relevance of these mechanisms and supports a causal interpretation of the results.

JEL Codes: E44, G32, Q54, L14.
Keywords: Transition risk, Cost of debt, Production networks, Environmental shocks, Default risk, Carbon emissions.

*Corresponding author: Alessandro Dario Lavia, alessandro.lavia@unito.it. Contact: Elisa Luciano, elisa.luciano@unito.it. We benefited from discussions with many including Peter Ireland, Jaromir Nosal, Rosen Valchev, Michele Boldrin, Ivan Petrella, Roberto Marfe, Ali Usman, Fabio Bagliano, Paolo Pin, Alexander Wagner, Steven Ongena, Carlos Ramirez, Giacomo Como, Fabio Fagnani, Marco Scarsini. We thank the seminar participants of HEC Paris, Politecnico of Milan and of the 6th Bank of Italy-LTI Conference.

1 Introduction

Firms are subject to climate physical risks, or losses due to natural disasters, and transition risks, due to regulation, reputation damages and technological obsolescence. Brown firms, usually identified with more polluting ones, are, by their nature, more exposed to transition risks. The question of how physical and transition risks, that we name environmental or climate risks, are transmitted into financial markets is central to both policy debates and academic research.

Transition risks have been found to affect significantly firms' profitability, as in Bartram and Sehoon (2022), and to be priced in equity markets, as shown by Bolton and Kacperczyk (2021). As concerns credit markets, there is growing evidence that firms with larger carbon footprints face higher borrowing costs. Ivanov et al. (2023) show that in the US borrowing conditions, including interest rates, are more severe for brown firms, following both the California cap-and-trade act and the Waxman-Markey act. Kacperczyk and Peydro' (2024) use a sample of global firms to show that, once banks commit to decarbonization, polluting firms receive less bank credit. A similar evidence has been detected by Delis et al. (2024), Altavilla et al. (2023), Ehlers et al. (2022), Reghezza et al. (2022). Evidence of an extra return on brown or conventional versus green corporate bonds has been provided by Zerbib (2019) and Baker et al. (2018), among others.

These findings have been taken as an indication that credit markets begin to internalize physical as well as transition risks. Yet the existing evidence remains essentially reduced-form and partial: it interprets climate risk as a firm-level attribute, mostly proxied by direct emissions, and studies its correlation with observed spreads. Such an approach neglects two essential aspects. First, firms do not operate in isolation, but are embedded in complex input-output networks, where shocks originating in one sector propagate to others. Second, interest rates are not exogenously given risk premia, but equilibrium outcomes that result from the interaction of production, financing, and default risk. A proper understanding of environmental risks propagation and financial market pricing requires a framework that makes these mechanisms explicit.

This paper is a first attempt to develop such a framework. Building on recent advances in the theory of production networks, we construct a static general equilibrium model with three types of agents: firms, consumers, and banks. Firms are heterogeneous in their environmental exposure. Green firms are subject only to physical risks, while brown firms are also exposed to transition risks. In the absence of shocks, brown firms may be more productive than green ones. Crucially, firms must decide ex ante on inputs and labor without being able to adjust to the realized state. This rigidity, reflecting real-world delays and irreversibilities in production and employment, generates the possibility of default when shocks realize, or ex post.

Shocks do not stop at the firm level. Because production uses intermediate goods, shocks propagate through input-output linkages. The Leontief inverse provides the natural tool to characterize this propagation: the primitive shocks to productivity are transformed into network-adjusted shocks that affect all downstream firms. Hence, a firm's effective exposure is determined not only by its own environmental risk but also by that of its suppliers.

Banks operate under perfect competition, are risk-neutral, and finance firms' liabilities. They are residual claimants, paid only after workers and suppliers, and are therefore directly exposed to default risk. Interest rates are set so that banks break even in expectation. This zero-profit condition delivers a simple but powerful characterization: the interest rate applied to any firm depends on the distribution of its total shocks, namely the primitive ones affecting its own output and the ones inherited from its suppliers. Productivity differences between green and brown firms do not directly affect borrowing costs; only risk does.

The model yields several implications. First, firm borrowing costs increase not only with direct exposure but also with the exposures of upstream suppliers. Thus even firms with low direct emissions may face high interest rates if they are embedded in brown supply chains, because primitive shocks to brown firms are greater than those to green firms. Second, interest rates on brown firms are weakly higher if the network is the same for brown an green firms, or if brown firms are more connected to brown and green to green. These results show how risk, network propagation, and endogenous pricing jointly determine the allocation of credit.

We bring the mechanisms highlighted by the model to the data by assembling a novel panel that combines U.S. firm-level balance sheet information and borrowing costs with sector-level carbon emissions from the Environmental Protection Agency and detailed input-output linkages from national accounts. This empirical design allows us to construct both a measure of direct carbon exposure and a model-consistent measure of network-adjusted emissions that captures firms' inherited exposure to transition risk through production networks. Our empirical analysis delivers two main findings. First, we show that exposure to carbon emissions is positively associated with firms' borrowing costs. When considered separately, both direct CO2 emissions and network-adjusted CO2 emissions are associated with significantly higher interest rates. This indicates that credit markets price climate-related risk and that firms operating in more emission-intensive environments face systematically higher financing costs.

Second, and more importantly, we show that both direct emissions and network-adjusted emissions matter once we isolate the residual component of the network measure. When direct CO_2 emissions are included together with the residualized network exposure, both coefficients remain economically large and statistically significant. This indicates that the network measure captures variation beyond firms' own emission levels.

These results suggest that banks do not price carbon risk solely based on the level of emissions in a given sector, but also account for firms' positions within production networks, in particular whether a firm is more strongly connected to carbon-intensive, or cleaner, upstream suppliers. Consistent with the model, borrowing costs reflect not only firms' direct emissions, but also the transition risk embedded in their supply chain exposure.

We further validate this mechanism by exploiting the 2016 entry into force of the Paris Agreement as an exogenous shift in climate policy expectations. Using difference-in-differences and event-study designs, we show that borrowing costs increase disproportionately for firms embedded in carbon-intensive production networks following the Agreement. The absence of differential pre-trends and the concentration of post-treatment effects in the network-based measure provide causal support for the interpretation that lenders revise credit pricing in response to anticipated transition risks transmitted through supply chains.

Taken together, the theory and evidence show that climate transition risk is priced in credit markets primarily through production networks. Borrowing costs reflect not only firms' own environmental footprint, but, more fundamentally, their exposure to upstream carbon risk and their position in the supply chain. The empirical results therefore strongly support the model's central prediction that transition risk operates as a systemic, network-amplified phenomenon rather than a purely firm-level attribute.

The remainder of the paper is structured as follows. Section 2 develops the theoretical model, introducing rigid production, stochastic shocks, and competitive banks, and derives the main results on how environmental risk and network position shape borrowing costs. In Section 3 we outline the empirical methodology and explain how we bring the model's predictions to the data, including the use of the Paris Agreement as a quasi-natural experiment. Section 4 concludes with a discussion of the broader implications for credit markets and the transition to a low-carbon economy.

Literature. The motivation of our paper is in the largely documented evidence on the different credit conditions applied to green and brown firms.

Empirically, a growing body of literature indeed shows that environmental performance affects corporate financing. In credit markets, Ivanov et al. (2023) show that in the US borrowing conditions from banks, including interest rates, are more severe for brown firms. This happens after both the California cap-and-trade act and the Waxman-Markey act. Kacperczyk and Peydro' (2024) use a sample of global firms to show that, once banks commit to decarbonization, polluting firms receive less credit. A similar evidence has been detected in the EU banking loans. Altavilla et al. (2023) demonstrate that euro-area banks charge higher interest rates to firms with greater carbon emissions and lower rates to firms committing to emissions reductions, even after controlling for default risk. These effects are more pronounced among banks publicly committed to decarbonization. Restrictive monetary policy amplifies both credit risk premia and emission-related pricing premia, tightening credit more for brown firms than for green ones. Related evidence from Ehlers et al. (2022) shows that polluting firms began paying a "carbon premium" on syndicated loans after the Paris Agreement. Also Delis et al. (2024) find that after the Paris Agreement, banks began pricing climate policy risk into syndicated loans, especially for firms with large fossil fuel reserves.

More broadly, earlier studies had already established that firms with poor CSR or environmental records face higher borrowing costs. Goss and Roberts (2011) show that socially irresponsible firms pay a premium on bank loans, while Chava (2014) document higher loan spreads for environmentally controversial firms. In bond markets, a modest "greenium" has been identified: Zerbib (2019) and Baker et al. (2018) find that green bonds carry slightly lower yields relative to comparable conventional bonds, consistent with investor preference channels highlighted by Pastor et al. (2021).

As a result, Gormsen and Huber (2023) show that since 2016 greener firms have enjoyed a structurally lower capital cost, roughly 1% lower than brown firms, reflecting both the equity and debt markets.

Face to these findings, up to our knowledge, there is no theoretical model which studies the role of transition shocks in determining the price of debt for green and brown sectors, in equilibrium. We claim to be the first in developing a model which incorporates this feature and thus provide a mechanism explaining how brown sectors end up being priced more than green ones, ceteris paribus. A key role in the explanation will be provided not only by the direct exposure of each firm to transition shocks, but also the exposure of its suppliers, which may suffer losses and reduce production because of transitions shocks, and transmit it downwards. We claim that the exposure of each firm to transition shocks therefore depends on its position and connectedness in the production network.

To formalize the role of shock transmissions thorough input-output chains, we draw on the literature on production networks, along the seminal contribution of Acemoglu et al. (2012). We build on the frameworks of Como et al. (2026) and Pellet and Tahbaz-Salehi (2023), who introduce rigidity in input contracting. Rigidity means that orders for intermediate goods and investments occur before shocks to production realize, and cannot be adjusted when the true state of the world becomes known, as we witnessed during the Covid-19 pandemic and recent disruptions from wars. Because of rigidities, firms suffer losses and may even default. Relative to those papers, here we distinguish green and brown firms based on the risks they suffer and the productivity they enjoy. These elements allow us to explain how sector-level environmental shocks propagate through supply chains and influence firm-specific borrowing costs, even for firms that do not directly increase emissions. In Pellet and Tahbaz-Salehi (2023) there was neither leverage nor green vs brown firms, in Como et al. (2026) there was leverage but no distinction between green and brown firms.

2 Theoretical Analysis

The theoretical model builds on Como et al. (2026). We develop a static general equilibrium in a production network with three types of agents: firms, consumers, and banks. All decisions are taken at time 0, prior to the realization of uncertainty. At time 1, the stochastic shocks on production are realized, production occurs, consumers consume, and payments between agents are settled. Uncertainty in production is introduced by defining a complete probability space $(\Omega, \mathcal{F}, P)$, on which random shocks are realized. Agents make ex-ante decisions at time 0, while outcomes and transfers materialize ex-post, at time 1.

2.1 Rigidity

Production processes in the real world are characterized by lags and irreversibilities. Orders of intermediate goods and investment in physical capital typically take time to be delivered, and therefore, even when there is uncertainty on the future state of the economy, are planned in advance, without the possibility of being tailored to the the state of the world that will realize. When a specific state occurs, these decisions are often irreversible. Similarly, labour is often hired in advance and cannot promptly be dismissed, at least in some jurisdictions. These forms of *rigidity*, which have recently been introduced in the network literature (Pellet and Tahbaz-Salehi (2023)), may generate default, if firms are levered. If, in a specific state, the costs of orders for intermediate goods and labour hired, as well as the interest rates to banks, are not covered by revenues, default occurs. To model all of this, we assume that the decision of firms about intermediate goods and labour are not state dependent. They are characterized by ex ante or *nominal* quantities, constant across possible states.

2.2 Firms

There are n firms or sectors in the economy[1], $\mathcal{V} = 1, 2, .., n$, partitioned into two groups. The first n_1 firms are classified as *green*, while the remaining $n - n_1$ firms are classified as *brown*. Brown firms are more risky than green, in that they can be hit by transition as well as physical shocks, but are more productive, all others equal. Transition shocks include in our view fines, penalties, and other reputational and regulatory risks, as well as technological obsolescence. Physical risks correspond to natural catastrophes and their consequences.

On the probability space $(\Omega, \mathcal{F}, P)$, we define two vectors of non-positive random variables: $\lambda : \Omega \to \mathbb{R}^n_-$ and $\epsilon : \Omega \to \mathbb{R}^n_-$, whose first n_1 components are equal to zero. These represent the primitive production log shocks, for brevity the *primitive production shocks*. Green firms are exposed only to λ, which captures physical risks. Brown firms are exposed to both λ and an additional transition shock ϵ. Let $\eta = \lambda + \epsilon$ denote the overall vector of primitive shocks. The previous assumptions on the components of η are equivalent to stating that shocks to brown firms are larger in absolute value than those to green firms, or that *green shocks dominate brown ones in the sense of first-order stochastic dominance.* Let G_k be the cumulative distribution function of $|\eta_k|$: for any $b > n_1$ (brown) and $g \leq n_1$ (green),

$$G_b(x) \leq G_g(x) \quad \text{for all } x \in \mathbb{R}^n, \text{ with strict inequality for some } x, \tag{1}$$

When considering η_k instead of its absolute value, since shocks are non-positive, the inequality reverses. If F_k denotes the cumulative distribution function of η_k:

$$F_b(x) \geq F_g(x), \quad \text{for all } x \in \mathbb{R}^n, \text{ with strict inequality for some } x, \tag{2}$$

[1]We use indifferently the term firm and sector because we assume that firms are perfect substitutes: each sector is made by homogeneous firms, and as one firm defaults it can be substituted by a perfectly equal one. This prepares the empirical part, since the US input-output matrix is available to us at the 6 digit sector and not at the individual firm level

The definition of transition and physical risks translates into the following assumption: $\eta_g \geq \eta_b$ in the sense of first-order stochastic dominance. Considering the absolute value of the shocks, $-\eta_g \leq -\eta_b$: later on, we label this saying that *green shocks are smaller than brown.*

To capture the basic trade-off between productivity and risk, and anticipating on the empirical evidence, we assume that brown firms, in the absence of shocks, are more productive than green ones. Let $Z \in \mathbb{R}^n_+$ denote the vector of total factor productivities. We assume Z_k is higher for brown firms ($k > n_1$) and lower for green firms ($k \leq n_1$), reflecting a higher baseline output among brown firms that compensates for their greater environmental exposure.

Firms use a Cobb-Douglas technology to produce output:

$$y_k^\eta = Z_k e^{\eta_k} l_k^{\beta_k} \prod_{j\in\mathcal{V}} (z_{jk}^\eta)^{A_{jk}}, \quad \text{for } k = 1, \ldots, n, \tag{3}$$

where:

- y_k^η is the output of firm k after the shock η,
- $Z_k > 0$ is the firm-specific productivity,
- l_k is labor employed,
- $\beta_k \geq 0$ is the labor share,
- z_{jk}^η is the input of good j used by firm k after the shock η,
- $A_{jk} \geq 0$ reflects the *importance* of input j in the production of firm k.

The case of no shocks is denoted by $\eta = 0$. We define $\alpha_k = \sum_{j\in\mathcal{V}} A_{jk}$ and assume constant returns to scale:

$$\alpha_k + \beta_k = 1. \tag{4}$$

Let c_k denote the consumption of good k.

We collect output, consumption, labor, and productivity in the vectors $y, c, l \in \mathbb{R}^n_+$, $Z \in \mathbb{R}^n_{++}$, and the input coefficients z_{jk} in a matrix $z \in \mathbb{R}^{n\times n}_+$. The production function of firm k can be compactly written as $y_k = F_k^{\eta_k}((z_{jk})_j, l)$, where $(z_{jk})_j$ is column k of the matrix z, which collects all the inputs to sector k.

The market clearing conditions are:

1. Goods market clearing:

$$y = z\mathbf{1} + c, \tag{5}$$

2. Labor market clearing:

$$l'\mathbf{1} = 1. \tag{6}$$

where $\mathbf{1}$ is a vector of ones.

We also collect the coefficients A_{jk} into a matrix $A \in \mathbb{R}_+^{n \times n}$. Under the mild assumption that the spectral radius of $A, \sigma(A)$ be smaller than one, the matrix A defines a Leontief inverse:

$$L = (I - A')^{-1}, \tag{7}$$

which has non-negative entries and plays a central role in shock propagation throughout the network, as we show now.

We assume that shocks propagate proportionally through the network: if a firm is hit by a shock that reduces output by a fraction p, it also reduces its supply of intermediate goods equally across suppliers and its supply of final goods by the same fraction. Formally, for all k, j such that $A_{kj} > 0$,

$$\frac{\tilde{y}_k}{y_k} = \frac{\tilde{z}_{kj}}{z_{kj}} = \frac{\tilde{c}_k}{c_k}. \tag{8}$$

where the values with a tilde are the actual or post-shock ones, while the ones without the tilde are the ex ante or nominal ones. Como et al. (2025) show that, under this assumption, the actual shocks received by firms are not the primitive shocks η, but the network-amplified shocks $\rho = L\eta$. A shock η_j originating in node j affects firm k if $L_{kj} > 0$:

$$\rho_k = \sum_{j \in \mathcal{V}} L_{kj} \eta_j. \tag{9}$$

The vector ρ has non-positive entries, which represent the total shocks to the different sectors, opposite to the primitive ones in η. This highlights the core network effect: production shocks can be magnified or attenuated depending on the structure of interfirm linkages. A sufficient condition for the total shocks to brown firms to be greater than the ones of green ones - as it happened for primitive ones - is that the Leontief matrix separates the two networks, or $L_{kh} > 0$ when both firms are either green or brown ($k \leq n_1, h \leq n_1$ and $k > n_1, h > n_1$), while $L_{kh} = 0$ in the mixed cases, when one is green and the other brown.

Lemma 1. *If $L_{kh} > 0$ for $k \le n_1, h \le n_1$ and $k > n_1, h > n_1$, $L_{kh} = 0$ otherwise, $\rho_g \ge \rho_b$ in the sense of first-order stochastic dominance.*

To complete the description of the firm side of the economy, we introduce

- w, unit cost of the employed labor (*wage*);
- p_k, unit price of the good produced by firm k.

Because of (8), the *actual* assets or revenues of firm k are

$$\mathcal{A}_k := p_k y_k^\eta = p_k e^{\rho_k} F_k^0((z_{jk})_j, l_k) \tag{10}$$

On the other hand, its *actual* liabilities due to intermediate goods and labor are

$$\mathcal{L}_k := \sum_{j \in \mathcal{V}} p_j z_{jk}^\eta + w l_k = \sum_{j \in \mathcal{V}} p_j e^{\rho_j} z_{jk} + w l_k. \tag{11}$$

Firms raise debt from financiers, which we call banks, As we explain below, the bank finances a fraction θ_k of the liabilities., so that θ_k represents the *leverage ratio* of firm k. The bank either receives the full payment of the loan and its interests, $(1 + r_k)\theta_k \mathcal{L}_k$, or gets only a recovery out of it. In the second case the firm looses the complement of the recovery with respect to the full payment, $(1 + r_k)\theta_k \mathcal{L}_k$, as default costs. So the firm always pays $(1 + r_k)\theta_k \mathcal{L}_k$. Then, firms' profits are:

$$\Pi_k((z_{jk})_j, l_k, \eta, p) := \mathcal{A}_k - (1 + r_k \theta_k)\mathcal{L}_k \tag{12}$$

They can take any sign: for the sake of simplicity, we name them profits also when they are negative.

Actual assets, liabilities, and profits can be expressed in terms of normalized quantities, as follows. Define the *normalized total shocks* as the total shocks which affect k, divided by their expectation under the measure P

$$\tau_k := \frac{e^{\rho_k}}{\mathbb{E}[e^{\rho_k}]}. \tag{13}$$

Define also the *normalized suppliers' shocks* as the normalized total shocks to the nodes who provide inputs to node k, weighted by their importance. Considering that log shocks to labor are zero by assumption and inputs to sector k include labor, with importance β_k, the suppliers' shocks are

$$\epsilon_k := \beta_k + \sum_{j\in\mathcal{V}} A_{jk}\tau_j. \tag{14}$$

Then, defining $s_k := \mathbb{E}[\mathcal{A}_k]$

$$\mathcal{A}_k = s_k\tau_k \tag{15}$$

$$\mathcal{L}_k = s_k\frac{\epsilon_k}{1+r_k\theta_k}. \tag{16}$$

2.3 Banks

Financiers form a continuum of risk-neutral banks operating under perfect competition. Each bank is randomly matched with a firm and provides financing to cover a fraction $\theta_k \in [0,1]$ of the firm's liabilities. Banks are paid only after firms have settled payments to workers and suppliers. Because of that, they are exposed to the possibility of default, which arises endogenously from the realization of shocks.

Let us assume that banks recover the minimum between the amount due $((1+r_k)\theta_k\mathcal{L}_k$ and the cash flow available to the firm after labour ad the other firms have been paid. This means that labour and providers of intermediate goods are senior to the bank, or that the former do not take any default risk, as stated, and that the bank recovers whatever is left after those are paid, up to its credit. The bank's profit from lending to firm k is:[2]

$$\mathcal{I}_k = \min(\max(\mathcal{A}_k - (1-\theta)\mathcal{L}, 0), (1+r)\theta_k\mathcal{L}_k) - \theta_k\mathcal{L}_k \tag{17}$$

$$= \min(\max(\mathcal{A}_k - \mathcal{L}, -\theta_k\mathcal{L}_k), r_k\theta_k\mathcal{L}_k). \tag{18}$$

Because banks operate competitively and are risk-neutral, they choose the interest rate r_k so as to break even in expectation. The zero-profit condition is:

$$\mathbb{E}[\mathcal{I}_k] = 0 \tag{19}$$

To simplify matters, in what follows we consider only the case in which liabilities are totally financed by debt, or $\theta = 1$. In that case, the condition that determines the interest rate for sector k becomes

[2]The payoff to bank k is the one described here if $\mathcal{A}_k \geq 0$, namely if its revenues are non-negative, i.e. the financed firm does not sell negative quantities at positive prices or positive quantities at negative prices. Since the only equilibrium features positive prices and quantities of goods, for every θ, the requirement is satisfied as a strict equality: $\mathcal{A}_k > 0$. The reader interested in the general case ($\mathcal{A}_k \geq 0$ as well as $\mathcal{A}_k < 0$), and in the proof that also this more general case leads to positive revenues, can look at Como et al. (2026).

$$\mathbb{E}\left[\min(\mathcal{A}_k - \mathcal{L}_k, r_k\mathcal{L}_k)\right] = \frac{s_k}{1+r_k}\mathbb{E}\left[\min(\tau_k(1+r_k) - \epsilon_k, r_k\epsilon_k\right]. \tag{20}$$

$$\mathbb{E}\left[\frac{-\epsilon}{1+r} + \min(\tau_k, \epsilon_k)\right] = 0. \tag{21}$$

and finally, since $\mathbb{E}[\epsilon_k] = 1$

$$\mathbb{E}[\min(\tau_k, \epsilon_k)] = \frac{1}{1+r_k}. \tag{22}$$

Therefore, the interest rate applied to each firm depends exclusively on its exposure to environmental shocks, and not directly on its productivity[3] Z_k, even tough brown firms are assumed to be more productive than green firms (i.e., $Z_i > Z_j$ for $i > n_1$ and $j \leq n_1$). Interest rates depend only on shocks, because default does. However, and consistently with the transmission of shocks from one firm to the other through the network, which may generate default cascades (see also Como et al. (2025)), the shocks that matter are not only the primitive ones, η, but the total ones, and specifically the total shocks to the sector and to its suppliers, τ and ϵ.

Interest rates are non negative, since $\mathbb{E}[\min(\tau_k, \epsilon_k)] \leq \mathbb{E}[\epsilon_k] = 1$.

2.4 Network effects on the cost of debt

In this section we study the relationship between risks and cost of debt. We show that, in equilibrium, the cost of debt on green firms is smaller than the one on brown ones in a number of cases. First, when the network linkages are uniform, in that all sectors have the same importance. Second, if the networks formed by green and brown firms are almost separated (brown firms are more important to brown, green to green) and third, if the two subnetworks are completely separated. Let us consider first the case in which labour - which is riskless - is not more important for green than for brown firms and all the importance coefficients - contained in the matrix A - are the same, $A_{jg} = A_{jb}$ for all indices j, g, b. This is an economically interesting case not only because all intermediate goods are equally important, be them produced or purchased by green or brown firms. But also because the transmission of the shocks in the network occurs in a uniform way. A matrix A of this type indeed implies that the off-diagonal terms of the Leontief matrix are the same: transmission of the primitive shocks to any other sector occurs in the same, *uniform* way. It is not difficult

[3] The results of the theoretical model are therefore robust with respect to the higher productivity assumption of brown firms. Interest rates of brown firms, under the conditions of the theorems below, would be greater than the green ones also in the absence of such an assumption.

to guess that, in such a "neutral" transmission environment, brown sectors, whose primitive shocks are more severe, but are affected in the same way as green sectors by the shocks of other sectors, remain more risky and therefore deserve a greater interest rate. Banks charge higher interest rates to brown firms in order to compensate for higher default risk. This is the content of the next Theorem, which is proven in Appendix A.

Theorem 1. *[Interest Rates with a uniform network] If green firms do not employ less labour than brown firms, i.e., $\beta_g \geq \beta_b$ for all $g \leq n_1$ and $b > n_1$, and $A_{jg} = A_{jb}$ for all indices j, g, b, then the interest rate applied to brown firms is weakly higher than that applied to green firms:*

$$r_b \geq r_g, \tag{23}$$

.

The theorem outcome results from the brown firms' higher exposure to both physical and transition risks, as a result of their direct - not network related - exposure.

The second circumstance we study is the one in which green firms not only are affected by less severe primitive shocks, but are also more exposed to shocks transmitted by other green firms, because the importance of green firms for them is higher than the one of brown firms. Both groups are equally exposed to brown shocks, because the total importance of brown sectors to them is the same. The intuition is that the networks are *almost separated*. The following theorem is also proved in Appendix A:

Theorem 2. *[Interest Rates with stronger linkages among brown vs green sectors] If green firms do not employ less labour than brown firms, i.e., $\beta_g \geq \beta_b$ for all $g \leq n_1$ and $b > n_1$, and green sectors are more important than brown ones in the production of green sectors, $\sum_{j \leq n_1} A_{gj} \geq \sum_{j \leq n_1} A_{bj}$ and $\sum_{j > n_1} A_{gj} = \sum_{j > n_1} A_{bj}$, $\beta_g \geq \beta_b$ for all indices $g \leq n_1, b > n_1$, then the interest rate applied to brown firms is weakly higher than that applied to green firms:*

$$r_b \geq r_g, \tag{24}$$

The extreme case of the previous network separation is the one in which the networks are *completely separated*, which is described by the hypotheses of Lemma 1. [4]. Based on the previous Theorem, it would give the minimum interest rate for green firms. The following Corollary is also proved in Appendix A:

[4]This is the circumstance which would be most welcome, at least in Europe, by the regulatory authority, in order to have very favourable loan conditions for green firms (see for instance EC (2024)). We briefly discuss this in the conclusions

Corollary 1. *[Interest Rates with separated networks] If green firms do not employ less labour than brown firms, i.e., $\beta_g \geq \beta_b$ for all $g \leq n_1$ and $b > n_1$, and the hypotheses of lemma 1 hold, then the interest rate applied to brown firms is weakly higher than that applied to green firms:*

$$r_b \geq r_g, \tag{25}$$

As a further corollary, it is easy to prove that, if there are no physical risks, green firms deserve a zero interest rate, because they run no default risk at all. Also in this case the proof is in Appendix A.

Corollary 2. *[Interest Rates with separated networks and no physical risks] If green firms do not employ less labour than brown firms, i.e., $\beta_g \geq \beta_b$ for all $g \leq n_1$ and $b > n_1$, the hypotheses of lemma 1 are satisfied and there are no physical risks ($\lambda = 0$), then the interest rate on green sectors is zero.*

All these results prove that risk-adjusted interest rates arise endogenously and reflect primitive environmental uncertainty, but also the way in which it is transmitted through the network.

2.5 Consumers

Consumers form a continuum and are modeled as a representative price-taking agent endowed with preferences over the n consumption goods. Preferences are described by a Cobb-Douglas utility function:

$$U(c) = \prod_{k \epsilon \nu} c_k^{\gamma_k}, \tag{26}$$

where $\gamma_k \geq 0$ is the consumer's weight on good k, and, in vector notation, $\gamma' \mathbf{1} = 1$. Because the coefficients of the utility function can also be smaller for green sectors than for brown ones, or can be all equal, preferences do not take into account any specific preference for greenness.

Like firms, the representative consumer makes a *rigid* nominal consumption plan at time 0, before the realization of shocks. Let $c = (c_k)_{k=1}^n$ denote this vector of nominal consumption choices. Once the shocks η are realized, because of (8), the actual consumption is scaled by the total shock affecting each good:

$$c_k^\eta = e^{\rho_k} c_k. \tag{27}$$

The consumer maximizes expected utility over the distribution of shocks:

$$\max_{c} \ \mathbb{E}[U(c^{\eta})] \tag{28}$$

subject to a state-dependent budget constraint. The consumer receives:

- wages w from labor,
- profits or losses Π_k from firms, which include default costs,
- and interest income $r_k\theta_k\mathcal{L}_k$ from lending via banks, since consumers own both firms and banks.

The total endowment is therefore:

$$E = w + \sum_{k\in\mathcal{V}} \Pi_k + \sum_{k\in\mathcal{V}} r_k\theta_k\mathcal{L}_k. \tag{29}$$

The consumer's budget constraint, accounting for the realized shock to consumption, is:

$$\sum_{k\in\mathcal{V}} e^{\rho_k} c_k p_k \leq E. \tag{30}$$

Finally, define the *(Bonacich) network centrality vector* as:

$$v = L'\gamma. \tag{31}$$

Each element v_k of this vector captures the relative importance of each sector in the network, weighted by consumer preferences and input-output linkages.

2.6 Theoretical Framework: Equilibrium

We now characterize the general equilibrium (GE) of the economy described above. Following Como et al. (2025), we can state that a rigid Walrasian equilibrium exists and is unique when banks competitively set interest rates to satisfy their zero-profit condition.

We first define it:

Definition 1 (Rigid Walrasian Equilibrium)**.** Consider a Cobb-Douglas economy $(\mathcal{V}, A, \beta, \gamma, Z)$ satisfying (4), with $\sigma(A) < 1$, equipped with an interest rate vector r and a leverage vector θ. Given a distribution P over primitive log-productivity shocks η, a *rigid Walrasian equilibrium* is a tuple (y, z, c, l, p, w) such that:

(i) For every firm $k \in \mathcal{V}$, the labor l_k and input bundle $(z_{jk})_j$ maximize expected profits $\mathbb{E}[\Pi_k]$ given prices p, wage w, interest rates r, and leverage θ.

(ii) The consumption bundle c maximizes expected utility $\mathbb{E}[U(c^\eta)]$, subject to the state-dependent budget constraint:

$$\sum_{k \in \mathcal{V}} e^{\rho_k} c_k p_k \leq E. \tag{32}$$

We also need some additional notation. First, define the cost of debt under continuous compounding, instead of discrete, as $\zeta_k = log(1 + r_k \theta_k)$. This reflects the cost of debt to the single firm k. Define the total cost of debt over the supply chain, as the sum of the costs over the chain, weighted by the corresponding Leontief coefficients:

$$\xi_k = \sum_{j \in V} L_{kj} \zeta_j \,. \tag{33}$$

Also, let us define the Leontief inverse obtained from the modified importance coefficients,

$$\bar{A}_{jk} := A_{jk} \tfrac{1}{1 + r_k \theta_k}, \quad \bar{L} := (I - \bar{A}')^{-1}, \tag{34}$$

and the corresponding Bonacich vector, together with the sum of its components weighted by the labour coefficients

$$\begin{aligned} \bar{\beta}_k &:= \beta_k \tfrac{1}{1 + r_k \theta_k} \\ v^\zeta &:= \tfrac{\bar{L}' \gamma}{\bar{\beta}' \bar{L}' \gamma}, \end{aligned} \tag{35}$$

The following result is adapted from Como et al. (2025):

Theorem 3 (Existence and Uniqueness)**.** *Consider a Cobb-Douglas economy* $(\mathcal{V}, A, \beta, \gamma, Z)$ *satisfying* (4)*. Then a unique rigid Walrasian equilibrium* (y, z, c, l, p, w) *exists, and satisfies:*

- ***nominal productions:***

$$y_k^0 = v_k^\zeta e^{-\xi_k} \,; \tag{36}$$

- ***nominal intermediate quantities:***

$$z_{jk}^0 = v_k^\zeta A_{jk} e^{-\zeta_k - \xi_j}; \tag{37}$$

- ***employed labor:***

$$l_k = v_k^{\zeta} \beta_k e^{-\zeta_k}; \tag{38}$$

- ***nominal household's consumption:***

$$c_k^0 = \frac{\gamma_k e^{-\xi_k}}{\bar{\beta}' \bar{L}' \gamma}; \tag{39}$$

- ***prices over wage:***

$$\frac{p_k}{w} = \frac{e^{\xi_k}}{\mathbb{E}[e^{\rho_k}]}; \tag{40}$$

So, the interest rates studied above are consistent with an equilibrium with environmental shocks.

3 Empirical Analysis

This section provides an empirical validation of the theoretical framework developed in Section 2. The model predicts that, in equilibrium, corporate borrowing costs reflect firms' exposure to climate transition risk, where such exposure extends beyond firms' own direct emissions and is transmitted through production networks. As a result, lenders price network adjusted carbon exposure rather than direct sectoral emissions alone.

We test these predictions using firm year data from Compustat combined with sector level emissions and input output information. Climate exposure is measured at the level of narrowly defined industries and assigned to firms operating within those sectors, allowing us to study how network embedded carbon intensity affects borrowing costs at the firm level.

Using a newly assembled panel covering the period 2011 to 2023, we document a robust positive relationship between borrowing costs and network adjusted emissions, even for firms operating in industries with relatively low direct emissions. Exploiting the entry into force of the Paris Agreement as an exogenous shift in climate policy expectations, we further provide causal evidence that credit markets internalize transition risk transmitted through production networks, in line with the model's predictions.

3.1 Dataset Construction

Our empirical analysis is conducted at the firm year level using financial information from Compustat. Firm level variables include balance sheet items, capital structure, profitability, liquidity, and borrowing cost measures. Environmental exposure, however, is constructed at the industry level. Emissions data are available at the six digit NAICS industry level, which represents the finest sectoral classification for which consistent emissions information can be combined with input output tables. Each firm year observation is therefore assigned the direct and network adjusted CO_2 exposure of its corresponding NAICS6 industry.

The network based measure of environmental exposure is constructed using national input output tables from the Bureau of Economic Analysis, which capture production linkages across industries. By combining sector level emissions with these input output relationships, we obtain a measure of carbon exposure that reflects both firms' own direct emissions and the emissions embodied in their upstream production networks. This approach allows us to capture transition risk transmitted along supply chains in a manner consistent with the theoretical framework.

Observations are included only if firms report a valid NAICS classification and non missing values for key financial variables, including interest expense, interest bearing debt, and total assets. To mitigate the influence of outliers driven by reporting issues or extreme accounting realizations, financial variables are trimmed at the top and bottom percentiles following standard practice. All trimming and winsorization procedures are applied consistently across variables and years.

Consistent with the focus of the theoretical model, we restrict the analysis to manufacturing industries, defined as two digit NAICS codes 31 to 33. Manufacturing sectors account for the majority of industrial emissions and exhibit substantial heterogeneity in both direct and network based carbon exposure. In contrast, service oriented industries display negligible direct emissions and limited upstream carbon intensity, making them less informative for identifying the role of production networks in the transmission of climate transition risk.

Cost of Debt Our main outcome variable is a firm level measure of the cost of debt, constructed as the ratio of interest expenses to interest bearing debt. For each firm i in year t, we define:

$$r_{i,t} = \frac{\mathtt{xint}_{it}}{\mathtt{dltt}_{it} + \mathtt{dlc}_{it}} \times 100, \tag{41}$$

where $\mathtt{xint}_{it}$ denotes interest expense and $\mathtt{dltt}_{it} + \mathtt{dlc}_{it}$ corresponds to long term and short term interest bearing debt.[5] This ratio captures the effective interest rate paid on all outstanding interest bearing obligations at the firm level, aggregating across maturities and

[5] We also build a measure based on total liabilities, as reported in Appendix B.2.1.

instruments.[6]

This firm level measure of the cost of debt is widely used in empirical corporate finance as a proxy for borrowing costs, not only when market yields are unavailable, but also when the objective is to capture the effective interest rate paid on a firm's entire stock of outstanding debt, as in Pittman and Fortin (2003). By construction, this measure aggregates interest expenses across short and long term obligations and relates them to the corresponding debt stocks, providing a comprehensive summary of firms' realized financing costs.

In our context, this choice is particularly appropriate because it aligns closely with the theoretical mechanism studied in the paper. In the model, interest rates are set by banks as equilibrium objects that reflect expected default risk on the full liability structure of the firm, rather than on a single debt instrument or maturity. The ratio of interest expenses to interest bearing debt therefore captures the relevant object priced by lenders when assessing transition risk as it materializes.

Moreover, focusing on the effective interest rate allows us to account for heterogeneity in debt maturity, contractual terms, and refinancing decisions that may themselves respond to climate policy expectations. Unlike quoted loan spreads or bond yields, which pertain to specific instruments at issuance, the realized cost of debt reflects the cumulative impact of climate related risk pricing across the firm's entire balance sheet. This makes it particularly well suited to studying how transition risk, both direct and network based, is internalized by credit markets in equilibrium.

Productivity Measure Firm level productivity is measured using a standard residual approach, following Solow (1957). Total factor productivity $Z_{i,t}$ for firm i at time t is computed as:

$$\mathrm{Z}_{i,t} = \log(\mathrm{Y}_{i,t}) - \left[\beta \log(\mathrm{H}_{i,t}) + (1-\beta) \log(\mathrm{K}_{i,t})\right], \tag{42}$$

where $\mathrm{Y}_{i,t}$ denotes firm level value added, $\mathrm{H}_{i,t}$ represents labor input, $\mathrm{K}_{i,t}$ denotes capital input, and β is the labor share of income. We use industry year specific labor and capital

[6]To mitigate the influence of implausible values driven by reporting issues or timing mismatches between interest expenses and debt stocks, we apply several data quality filters. We drop observations with missing NAICS codes or missing values for interest expense or interest-bearing debt. We winsorize key balance sheet variables, including `xint`, `dltt`, `dlc`, `lt`, and `at`, at the 2.5th and 97.5th percentiles. Based on these cleaned components, we compute the implied interest rate and initially restrict the sample to observations with interest rates less than or equal to 30 percent. We then winsorize the resulting interest rate measure at the 5th and 95th percentiles, following the approach in Pittman and Fortin (2003). The same 30 percent cutoff and 5 percent winsorization are applied to the alternative interest rate measure based on total liabilities. Finally, we winsorize leverage at the 2.5th and 97.5th percentiles to reduce the influence of extreme capital structure realizations.

cost shares from the NBER CES database to construct this measure, following Becker et al. (2021). Under the assumptions of constant returns to scale, cost minimizing firms, and competitive factor markets, observed factor cost shares equal output elasticities, so that the Solow residual provides a valid measure of revenue total factor productivity.

Under these assumptions, productivity can be recovered directly from observed value added and factor inputs without estimating a production function, which avoids potential biases associated with simultaneity and functional form misspecification.

Network-Adjusted Emissions and Transition Risk Propagation. A central contribution of this paper is the construction of a network-based measure of carbon exposure that captures how transition risk propagates through production linkages across industries. To this end, we distinguish between direct emissions, also known as *scope 1*, and our novel measure of network-adjusted emissions, both defined at the six-digit NAICS level. While direct CO_2 emissions measure the carbon directly generated by a sector's own production activities, they fail to account for the substantial indirect exposure arising from reliance on upstream inputs produced in carbon-intensive industries.

Our network-adjusted emissions measure explicitly incorporates inter-industry dependencies and weights upstream emissions by their economic relevance in each sector's production structure, in line with the theoretical framework developed in this paper. This approach recognizes that sectors may face heightened transition risk even in the absence of their own direct emissions, simply due to their position within carbon-intensive supply chains.

We construct network-adjusted CO_2 emissions sector-by-sector by combining observed six-digit NAICS sectoral direct emissions with input-output relationships across industries. Let $\mathtt{CO2}_t$ denote the vector of observed sectoral direct emissions at time t, and let $\mathbf{L}_t$ denote the corresponding Leontief inverse matrix capturing both direct and indirect production linkages across sectors. Consistent with the calibration of the model [7], network-adjusted emissions for sector k are then computed as:

$$\mathtt{Network\ CO2}_{k,t} = \sum_{j \in \nu} \mathbf{L}_{jk,t} \cdot \mathtt{CO2}_{j,t} = \sum_{j \in \nu} \sum_{h=0}^{+\infty} (\mathbf{A}'_{jk})^h \cdot \mathtt{CO2}_{j,t} \qquad \forall i \in k \tag{43}$$

This construction propagates direct emissions throughout the entire production network, weighting upstream carbon intensity by each sector's dependence on intermediate inputs while allowing financial conditions to amplify exposure. As a result, network-adjusted CO_2

[7] In Appendix B.1.6 we show that the $\mathbf{A}$ matrix need to be scaled by by $(1 + r_t)$, where r_t represents the sector-specific output interest rate.

Figure 1: Histograms of `CO2` and `Network CO2` Emissions

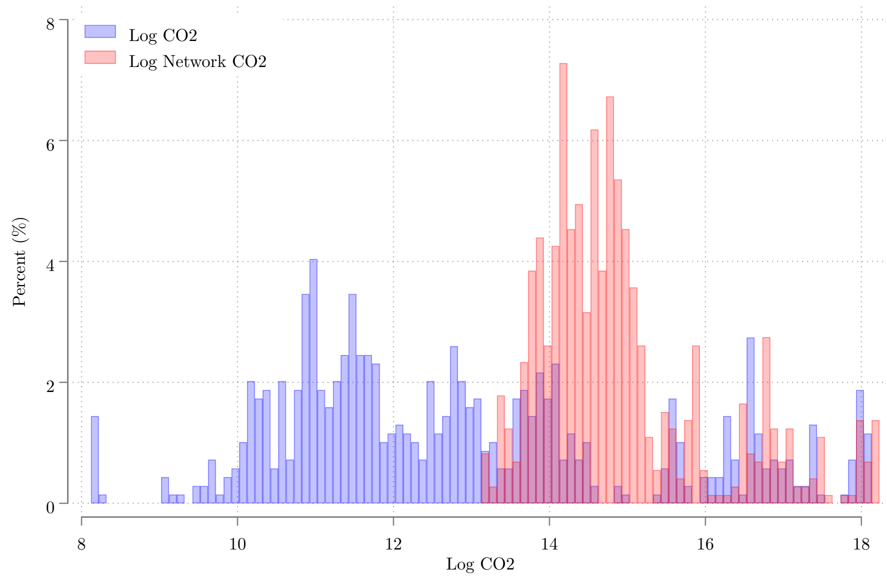

Notes: This figure displays histograms of firm-level `CO2` (blue bins) emissions and `Network CO2` (red bins) emissions, both expressed in logarithmic scale. Each histogram is based on approximately 6,000 observations. The y-axis reports the relative frequency of observations, expressed as percentages. `CO2` measures direct emissions at the firm level, while `Network CO2` captures emissions associated with firms' production networks as defined in this section.

captures the total embodied emissions associated with each sector's output, reflecting both direct emissions and those inherited through supply-chain linkages.

By explicitly accounting for the structure of production networks, this measure provides a more accurate representation of firms' effective carbon exposure and the channels through which transition risk is transmitted across industries. It is particularly relevant in carbon-intensive economies, where upstream shocks may generate substantial downstream financial consequences.

Figure 1 illustrates the distribution of both `CO2` and `Network CO2` in log scale, where unit of measure is expressed in million metric tons of CO_2 equivalent ($MtCO_2e$), calculated using 100-year global warming potentials from the IPCC Fourth Assessment Report (AR4). As discussed above, restricting the analysis to the manufacturing sector allows us to isolate emissions dynamics in sectors with substantial and economically relevant carbon intensity. The distribution of direct emissions is strongly right-skewed, with a large concentration of sectors exhibiting low emission levels and a long right tail driven by a small number of highly polluting industries.

In contrast, the network-adjusted measure substantially reweights emissions across sectors by accounting for input–output linkages. Sectors that appear to generate negligible direct emissions exhibit positive network-adjusted emissions due to their reliance on carbon-intensive upstream suppliers. This redistribution highlights how embodied emissions propagate through production networks, effectively assigning carbon exposure to downstream sectors according to their economic dependence on polluting industries.

3.2 Descriptive Statistics

Table 1 reports descriptive statistics for the main variables used in the empirical analysis. The sample covers the period from 2013 to 2020, providing balanced coverage around the entry into force of the Paris Agreement, which will be part of this empirical analysis. All financial variables are measured at the firm-year level, while emissions variables are constructed at the six-digit NAICS industry level and assigned to firms operating within those sectors.

The firm-level cost of debt, measured as interest expenses relative to interest-bearing debt, has an average of 6.18 percent and a standard deviation of 5.07 percentage points. The distribution exhibits moderate right skewness, with values ranging from close to zero to just above 23 percent. As discussed in Section 3, extreme values of interest rates are trimmed to mitigate the influence of reporting issues and timing mismatches between interest expenses and debt stocks.

Carbon emissions display substantial heterogeneity across industries. Direct emissions, measured by `CO2` and expressed in million metric tons of CO_2 equivalent ($MtCO_2e$), have an average of 4.26 and a standard deviation of 16.96, with a highly right-skewed distribution and a maximum exceeding 182 $MtCO_2e$. Network-adjusted emissions exhibit a higher average of 6.35 and similarly large dispersion, with even stronger right skewness. This reflects the amplification of carbon exposure through upstream production linkages, whereby sectors with low or moderate direct emissions inherit substantial embodied emissions from carbon-intensive suppliers.

The number of observations for emissions variables is lower than for financial variables because emissions are observed at the sector-year level rather than the firm-year level. In particular, log emissions are available for 636 sector-year observations for direct emissions and 658 for network-adjusted emissions, reflecting the number of six-digit NAICS manufacturing sectors with non-missing emissions data in each year of the sample period. These sector-level measures are then mapped to firms operating within the corresponding industries.

Table 1: Descriptive Statistics

	Average	SD	Min	Max	Skewness	Count
Interest Rate (Debt)	6.18	5.07	0.00	23.03	0.96	6965
`CO2`	4.26	16.96	0.00	182.51	7.59	6986
`Network CO2`	6.35	17.85	0.45	190.91	7.38	6982
Log Z	0.40	1.12	-6.35	2.18	-1.15	1579

Notes: This table reports descriptive statistics for firm–year observations. Borrowing costs are measured as interest expenses relative to interest-bearing debt or total liabilities. Direct emissions (`CO2`) and network-adjusted emissions (`Network CO2`) are constructed at the six-digit NAICS industry level and expressed in million metric tons of CO_2 equivalent ($MtCO_2e$), using 100-year global warming potentials from the IPCC Fourth Assessment Report (AR4). Emissions measures are assigned to firms based on industry affiliation.

Firm-level productivity, measured by the logarithm of total factor productivity, $\log \mathrm{Z}_{i,t}$, exhibits considerable dispersion, with an average of 0.40 and a standard deviation of 1.12. The distribution is markedly left-skewed, with a minimum of -6.35, indicating the presence of a subset of firm-year observations characterized by persistently low productivity levels. At the same time, the upper tail reaches values above 2, reflecting substantial heterogeneity in productive performance across firms.

Overall, the descriptive statistics highlight pronounced variation in borrowing costs, environmental exposure, and productivity across firms and sectors. Emissions are highly concentrated among a small subset of industries, while network-adjusted emissions redistribute carbon exposure across downstream sectors. These patterns motivate the subsequent analysis, which examines how both direct and network-based carbon exposure shape the pricing of corporate debt.

Classification of Brown and Green Sectors To compare outcomes across environmentally intensive and less intensive industries, we classify sectors as *brown* or *green* based on their relative carbon exposure within each year. Specifically, for every year, we consider the cross-industry distribution of lagged log carbon emissions and classify a sector as emission intensive (brown) if its lagged log emissions lie at or above the 80th percentile of the yearly distribution, while sectors below this threshold are classified as green. Using a high-emission cutoff rather than the median ensures that the brown group captures industries with clearly elevated environmental exposure, thereby isolating the upper tail of the carbon distribution where transition risk is most likely to be economically salient. Using lagged emissions further guarantees that sector classification is predetermined with respect to contemporaneous

Table 2: `Network CO2` Emissions and Productivity

	Average	SD	Min	Max	Skewness	Count
Green	0.37	1.11	-6.35	2.18	-1.13	1371
Brown	0.63	1.14	-4.60	2.16	-1.36	208
Total	0.40	1.12	-6.35	2.18	-1.15	1579

Notes: This table reports summary statistics for firm-level productivity across green and brown manufacturing sectors. Sector classification is based on network-adjusted carbon emissions measured at the six-digit NAICS industry level. In each year t, a sector is defined as brown if its lagged log `Network CO2` emissions are at or above the 80th percentile of the cross-industry distribution, and as green otherwise. Productivity is measured as the log of total factor productivity, $Z_{i,t}$, constructed using a standard Solow residual.

financial outcomes.

This classification is implemented separately for direct emissions measured at the industry level and for network-adjusted emissions that capture exposure through input-output linkages. The resulting indicators therefore distinguish brown and green sectors either on the basis of their own lagged emissions or on the basis of their inherited exposure through production networks. Focusing on the upper tail of the distribution sharpens the contrast between high- and low-exposure sectors and reduces the risk that marginal emitters dilute the estimated differences.

Because the classification is time varying and constructed relative to the within-year cross-industry distribution, it allows for consistent comparisons of financial outcomes across sectors with differing degrees of environmental risk while abstracting from aggregate trends in emissions over time. In the preliminary descriptive statistics, however, we summarize outcomes by averaging over the full sample period to provide a first, unconditional comparison between brown and green sectors.

Table 2 reports summary statistics for productivity across green and brown sectors. Green sectors have an average log productivity of 0.37, with a standard deviation of 1.11, a minimum value of -6.35, and a maximum of 2.18. Brown sectors exhibit a higher mean productivity of 0.63 and a slightly larger dispersion, with a standard deviation of 1.14, a minimum of -4.60, and a maximum of 2.16. The brown group represents the upper tail of the emissions distribution and comprises 208 observations, compared with 1,371 observations for green sectors, for a total sample size of 1,579 observations.

Both groups display negatively skewed productivity distributions, with skewness equal to

-1.13 for green sectors and -1.36 for brown sectors, indicating exposure to low productivity realizations. The skewness is slightly more pronounced among brown sectors, suggesting a thicker lower tail despite their higher mean productivity. At the same time, the upper tails of the distributions are similar across the two groups, with nearly identical maximum values.

Overall, the descriptive evidence indicates that sectors located in the upper tail of the carbon exposure distribution combine higher average productivity with substantial dispersion in outcomes. These patterns motivate the subsequent empirical analysis investigating how environmental exposure influences credit market conditions and propagates through production networks.

Consistent with our theoretical framework, which allows for higher productivity in emission-intensive sectors, we provide this descriptive statistics table to highlight empirical patterns that are well grounded in the existing literature. In particular, the observed productivity advantage of high-emission industries aligns with the argument that environmental constraints interact with firms' technological choices and competitive strategies, as emphasized by Porter (1991) and Porter and van der Linde (1995), who stress that productivity and competitiveness need not be adversely affected by environmental regulation. Moreover, our findings are consistent with evidence from the European context showing that environmental policies can coexist with, and in some cases even enhance, economic performance and productivity outcomes, as documented by Dechezlepretre et al. (2021). Taken together, these descriptive statistics validate our modeling assumptions and motivate the subsequent empirical analysis of how environmental exposure shapes credit market conditions and propagates through production networks.

3.3 Direct vs Network-Adjusted Emissions

We examine how exposure to direct and network-based carbon emissions affects corporate borrowing costs using a panel regression framework with high-dimensional fixed effects. The dependent variable is the firm-level effective interest rate paid by firm i in year t. Carbon exposure variables are constructed at the six-digit NAICS industry level and assigned to firms operating within those sectors. The explanatory variables are expressed in logarithms, so coefficients measure semi-elasticities: they capture the percentage-point change in the interest rate associated with a one percent change in carbon exposure.

Our baseline specification relates firm-level borrowing costs to lagged industry-level emis-

Table 3: Direct vs Network-Adjusted Emissions

	Interest Rate (1)	Interest Rate (2)	Interest Rate (3)
Log $\texttt{CO2}_{\text{t-1}}$	2.434*		2.792**
	(0.910)		(0.952)
Log $\texttt{Network CO2}_{\text{t-1}}$		2.546**	
		(0.890)	
Log $\texttt{Network CO2}_{\text{t-1}}$			1.922*
			(0.883)
Firm FEs	Yes	Yes	Yes
Year Fes	Yes	Yes	Yes
Sectors SEs	Yes	Yes	Yes
R-squared	0.686	0.682	0.686
N	1744	1778	1744

Notes: This table reports estimates from the baseline regression model in Equation 44. The dependent variable is the firm-level interest rate. Column (1) includes lagged direct industry-level `CO2` emissions. Column (2) replaces direct emissions with lagged network-adjusted emissions (`Network CO2`). Column (3) includes lagged direct emissions and the residual component of network-adjusted emissions orthogonalized with respect to direct emissions. All specifications include log leverage and log firm size as controls, as well as firm and year fixed effects. Standard errors are clustered at the firm level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

sions, with $\mathcal{C} = \{\texttt{CO2}, \texttt{Network CO2}\}$, and takes the form

$$r_{i,t} = \alpha + \beta_1 \log \mathcal{C}_{k(i),t-1} + \beta_2 X_{i,t-1} + \gamma_i + \delta_t + \varepsilon_{i,t}, \tag{44}$$

where $k(i)$ denotes the six-digit NAICS industry in which firm i operates. The vector $X_{i,t-1}$ includes firm-level controls that affect borrowing costs, specifically log leverage and log firm size. All specifications absorb firm fixed effects γ_i, which control for time-invariant firm characteristics, and year fixed effects δ_t, which capture aggregate macroeconomic conditions, credit market cycles, and common regulatory developments. Standard errors are clustered at the firm level to account for serial correlation within firms over time. For brevity, the coefficients on control variables and fixed effects are included in the regressions but not reported in the table.

Table 3 reports the main regression results across three specifications. Column (1) relates firm-level borrowing costs to lagged direct emissions at the industry level. The coefficient on log `CO2` is 2.434 and statistically significant at the 10 percent level. This estimate implies that a one percent increase in industry-level direct emissions is associated with an increase of approximately 2.4 basis points in firms' borrowing costs, conditional on firm characteristics and aggregate time effects.

Column (2) replaces direct emissions with the network-adjusted emissions measure. The estimated coefficient on log `Network CO2` is 2.546 and statistically significant at the 5 percent level. This result indicates that firms operating in sectors embedded in more carbon-intensive production networks face significantly higher borrowing costs, even after controlling for leverage, size, firm fixed effects, and year fixed effects.

Column (3) includes lagged direct emissions together with the residual component of network-adjusted emissions obtained by orthogonalizing `Network CO2` with respect to direct `CO2`. This residualized measure isolates variation in network exposure that is not mechanically driven by a sector's own emissions. In this specification, the coefficient on log `CO2` increases to 2.792 and remains statistically significant at the 5 percent level, while the coefficient on the residualized network measure is 1.922 and statistically significant at the 10 percent level.

These findings indicate that both direct emissions and the network-based component of carbon exposure are priced in credit markets. Importantly, the residualized specification demonstrates that network exposure exerts an independent effect on borrowing costs beyond firms' own direct emissions. The magnitude of the network coefficient remains economically meaningful, suggesting that creditors internalize transition risk transmitted through supply-chain linkages rather than focusing exclusively on firms' direct carbon footprint.

Taken together, the results in Table 3 highlight the importance of accounting for network-based carbon exposure when assessing the pricing of corporate credit. Once leverage, firm size, firm heterogeneity, and aggregate time effects are fully absorbed through controls and fixed effects, exposure to carbon-intensive production networks continues to exert a robust and economically meaningful impact on borrowing costs. These findings are consistent with a view of climate transition risk as a systemic phenomenon embedded in production networks.

3.4 Causality and the Paris Agreement: A Network Perspective

In this subsection, we move beyond correlations and investigate the causal impact of climate policy expectations on firms' borrowing costs. Our analysis focuses specifically on the network-adjusted emissions measure, which constitutes the central object of interest in this study. In the previous sections, we document that firms' exposure to carbon-intensive production networks plays a statistically and economically significant role in explaining the cross-sectional variation in interest rates charged by banks. Having established the relevance of the network channel, we now examine whether this relationship reflects a causal effect that becomes particularly salient in the presence of a major climate policy shock.

To identify causality, we exploit the entry into force of the Paris Agreement in late 2016 as a global, externally imposed shock to climate governance. The Paris Agreement marked a coordinated international commitment to long-run decarbonization and signaled a structural shift in expectations regarding the future stringency of environmental regulation across countries and industries. Crucially, this shift was broad-based and externally determined, rather than driven by firm-level financial conditions or sector-specific credit dynamics.

Importantly for identification, the Paris framework applied to all major emitting economies, including the United States. Although adopted in December 2015, the agreement formally entered into force in November 2016 following ratification by a critical mass of large emitters. Under the Obama administration, the United States formally joined the agreement, signaling a credible commitment to nationally determined contributions and to a trajectory of progressively tighter climate regulation. While the Trump administration announced its intention to withdraw in June 2017, institutional treaty rules implied that the United States remained a party until November 2020. The subsequent reentry under the Biden administration in early 2021 further reinforced the perception that international climate policy reflects a persistent and structural shift rather than a transitory political stance. Our sample period spans this entire phase of activation and implementation, making the Paris Agreement a suitable quasi-natural experiment.

We exploit this institutional setting using two complementary empirical strategies tailored to the network exposure measure. First, we implement a difference-in-differences design comparing the evolution of borrowing costs between firms more exposed to carbon-intensive production networks ("network-brown") and those with lower network exposure ("network-green") before and after the Paris Agreement. Second, we complement this analysis with an event-study approach that traces the dynamic adjustment of credit markets around the policy event and allows us to formally assess the parallel trends assumption underlying the difference-in-differences framework.

Focusing on the network-adjusted measure is conceptually central to our research question. If financial markets internalize transition risk not only through firms' own emissions but also through their embeddedness in carbon-intensive supply chains, then the repricing of credit risk should materialize most clearly along the network dimension following a credible climate policy shock. The Paris Agreement provides precisely such a shock: it generated a common and plausibly exogenous shift in expectations about future regulatory stringency, thereby increasing the salience of indirect carbon exposure transmitted through production networks.[8]

By leveraging this quasi-natural experiment, we directly assess whether firms' exposure to carbon-intensive production networks causally affects their cost of debt financing. In doing so, we isolate the network channel as the key mechanism through which climate transition risk is transmitted to credit markets.

Difference-in-Differences Approach (Network Exposure). We now implement a difference-in-differences design that exploits cross-industry heterogeneity in *network-adjusted* carbon exposure and compares firm-level borrowing costs before and after the entry into force of the Paris Agreement. Consistent with the core objective of this study, treatment status is defined exclusively using the network-based emissions measure.

Industries are classified as network-brown if their lagged growth in log network-adjusted emissions lies above the 80th percentile of the cross-industry distribution within a given year. In other words, the treatment indicator equals one for industries in the top 20 percent of the distribution of network carbon growth and zero otherwise. This high-exposure cutoff ensures that the treated group captures sectors with pronounced increases in inherited carbon intensity, thereby isolating industries most likely to face elevated transition risk. Firms operating in these network-brown industries constitute the treated group, while firms in network-green industries form the control group.

The identifying assumption is that, absent the Paris Agreement, borrowing costs for firms in network-brown and network-green industries would have followed parallel trends. The Paris Agreement represents a common and externally imposed policy shock, and identification comes from differential exposure to transition risk through production networks.

We estimate the following two-way fixed effects specification:

$$r_{i,t} = \alpha + \delta \left(\texttt{Network B}_{k(i)} \times \text{Post}_t \right) + \mu_i + \lambda_t + \varepsilon_{i,t}, \tag{45}$$

where $r_{i,t}$ denotes the firm-level interest rate, $\texttt{Network B}_{k(i)}$ indicates whether firm i operates in a network-brown industry (top 20 percent of network emission growth), and Post_t equals

[8]A growing literature exploits the entry into force of the Paris Agreement as a quasi-natural experiment to identify the causal effects of climate policy expectations on economic and financial outcomes. Prior work documents improvements in global environmental efficiency using a fuzzy regression discontinuity design (Salman et al., 2022) and positive stock market responses for firms exposed to green revenue opportunities (Kruse et al., 2024). These findings support the interpretation of the Paris Agreement as an externally imposed shift in expectations about future climate policy.

one for years 2017 and later.

Firm fixed effects μ_i control for time-invariant firm characteristics such as baseline credit quality, business model, or managerial practices. Year fixed effects λ_t absorb aggregate macroeconomic conditions, credit supply shocks, and common regulatory developments affecting all firms. Including both firm and year fixed effects is essential in a multi-year difference-in-differences design. Firm fixed effects remove permanent cross-sectional differences between treated and control firms, while year fixed effects absorb common time shocks. Identification therefore relies solely on within-firm changes in borrowing costs relative to the contemporaneous evolution of the control group. Standard errors are clustered at the firm level to account for serial correlation.

Importantly, we do not include additional firm-level controls in the baseline DiD specification. Because the Paris Agreement constitutes an externally imposed policy shock, and because firm and year fixed effects already absorb time-invariant heterogeneity and aggregate trends, the interaction term isolates the differential post-2016 change in borrowing costs attributable to network-based carbon exposure.[9]

Table 4 reports the estimates. Column (1), which includes firm fixed effects but not year fixed effects, yields a coefficient of 0.868, statistically significant at the 1 percent level. Column (2) adds year fixed effects, thereby controlling for aggregate credit market conditions and common macroeconomic shocks. The coefficient remains positive and statistically significant, with an estimate of 0.795. The stability of the estimate across specifications indicates that the differential increase in borrowing costs is not driven by time-invariant firm characteristics or common time effects, but instead reflects a relative repricing of firms exposed to carbon-intensive production networks.

Overall, the difference-in-differences results provide causal evidence that the Paris Agreement triggered a differential increase in borrowing costs for firms embedded in carbon-intensive supply chains. Consistent with the theoretical framework, credit markets internalize transition risk primarily through network-based exposure, particularly for industries located in the upper tail of the network carbon growth distribution.

Event Study Approach. To complement the difference in differences analysis and to examine the dynamic adjustment of borrowing costs around the Paris Agreement, we estimate

[9] Prior to constructing the treatment indicator, we winsorize emissions and log emissions at the 1st and 99th percentiles and trim extreme values to mitigate the influence of outliers. Network-brown status is defined using the 80th percentile of the annual cross-industry distribution of lagged log changes in network-adjusted emissions, consistent with the empirical construction described in the data section.

Table 4: Difference-in-Differences Estimates: Paris Agreement and Network Exposure

	Interest Rate (1)	Interest Rate (2)
`Network B` x T	0.868***	0.795**
	(0.261)	(0.280)
Firm FEs	Yes	Yes
Year FEs	No	Yes
Sectors SEs	Yes	Yes
R-squared	0.597	0.607
N	3250	3250

Notes: This table reports difference-in-differences estimates of the effect of the Paris Agreement on firm-level borrowing costs. Treatment status is defined using network-adjusted carbon exposure at the six-digit NAICS level. Industries are classified as brown if their lagged growth in log network-adjusted emissions lies above the 80th percentile of the annual cross-industry distribution. The post indicator equals one for years 2017 and later. Column (1) includes firm fixed effects only. Column (2) includes both firm and year fixed effects. Standard errors are clustered at the firm level and reported in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels.

an event study specification centered on the policy implementation year 2017, with 2016 serving as the omitted baseline year. While the DiD framework identifies an average post treatment effect, the event study design allows us to assess the credibility of the parallel trends assumption and to trace the timing and persistence of credit market responses to climate policy expectations.

Treatment status is defined using the same network based classification employed in the DiD analysis. Industries are classified as network brown if their lagged growth in log network adjusted emissions lies at or above the 80th percentile of the annual cross industry distribution. Firms operating in these industries are assigned a treatment indicator equal to one, while firms in the remaining industries form the network green control group. By focusing on the top 20 percent of the distribution, this classification isolates industries experiencing pronounced increases in inherited carbon exposure, thereby concentrating on sectors most plausibly exposed to elevated transition risk following a credible climate policy shock.

We estimate the following two way fixed effects event study specification:

$$r_{i,t} = \alpha + \sum_{s \neq -1} \beta_s \left(\texttt{Network B}_{k(i)} \times 1\{t - 2017 = s\} \right) + \gamma_i + \lambda_t + \varepsilon_{i,t}, \tag{46}$$

where $r_{i,t}$ denotes the firm level effective interest rate, $\texttt{Network B}_{k(i)}$ is the network brown indicator for firm i's six digit NAICS industry $k(i)$, and $1\{t - 2017 = s\}$ is an event time dummy equal to one in relative year s. The year immediately preceding implementation, $s = -1$, corresponding to 2016, is omitted from the regression, so each coefficient β_s measures the differential borrowing cost between network brown and network green firms in relative year s compared to the pre Agreement baseline.

Firm fixed effects γ_i absorb all time invariant firm characteristics, including persistent differences in credit quality, capital structure, and business models. Year fixed effects λ_t control for aggregate macroeconomic conditions, common credit supply shocks, monetary policy changes, and other regulatory developments affecting all firms simultaneously. Standard errors are clustered at the firm level to allow for serial correlation in borrowing costs over time.

The pre treatment coefficients, that is for $s < 0$, provide a direct test of the parallel trends assumption. Under the identifying assumption, borrowing costs for network brown and network green firms should evolve similarly prior to the Paris Agreement, implying that these coefficients should be close to zero and statistically insignificant. Post treatment coefficients, that is for $s \geq 0$, capture the dynamic repricing of transition risk following the policy shock and indicate how rapidly and persistently credit markets incorporate network based exposure into borrowing costs.

The estimated coefficients are plotted against event time and normalized to zero in $t = -1$. Confidence intervals correspond to 90 percent intervals, constructed as plus or minus 1.64 times the firm clustered standard errors. The vertical reference line marks the implementation year, and the horizontal line at zero corresponds to no differential borrowing cost relative to the baseline period.

The pre treatment coefficients are small in magnitude and statistically indistinguishable from zero, providing no evidence of differential pre trends between network brown and network green firms. This absence of pre trends supports the credibility of the identifying assumption underlying both the DiD and event study specifications.

Following the Paris Agreement, borrowing costs for network brown firms increase relative to the control group. The differential effect strengthens in the first and second years after implementation, with peak increases approaching approximately one percentage point. The estimated effects remain positive over subsequent years, suggesting a persistent repricing of firms embedded in carbon intensive production networks. Although confidence intervals widen at longer horizons, reflecting reduced precision, the overall pattern is consistent with

Figure 2: Event Study Estimates Around the Paris Agreement – Network Exposure

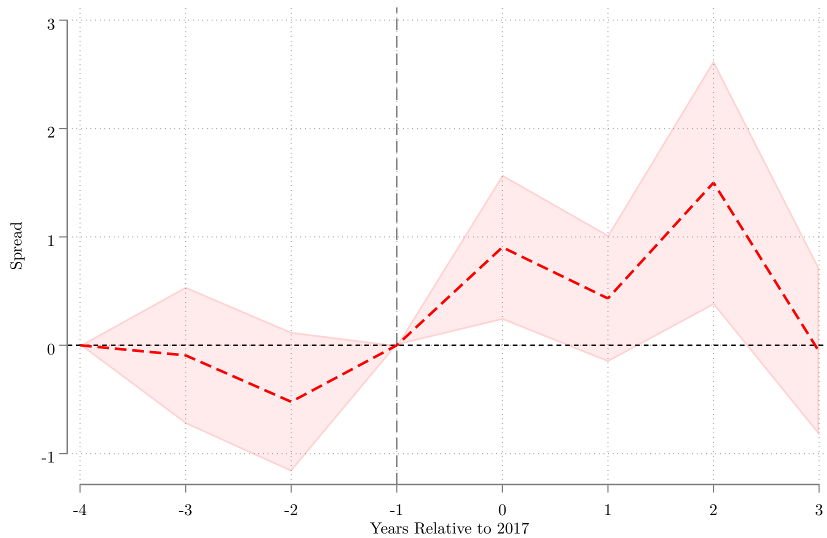

Notes: This figure plots event-study coefficients β_s from regressions of firm-level borrowing costs on event-time dummies interacted with a network-brown indicator. Event time is measured relative to 2017, and coefficients are normalized to $t = -1$ (2016). Industries are classified as network-brown if their lagged growth in log network-adjusted emissions lies at or above the 80th percentile of the annual cross-industry distribution. The solid red line reports point estimates; the shaded area shows 90 percent confidence intervals based on firm-clustered standard errors. A vertical dashed line marks the year before the implementation of the Paris Agreement.

a sustained adjustment rather than a temporary response.

Taken together, the event study evidence reinforces the DiD results. The absence of statistically significant pre treatment differentials combined with a clear post 2017 divergence in borrowing costs supports a causal interpretation. The Paris Agreement triggered a relative increase in borrowing costs for firms positioned within carbon intensive production networks, consistent with the view that transition risk is transmitted through supply chain linkages rather than solely through firms' direct emissions.

Summary. Taken together, the difference-in-differences and event-study analyses provide consistent causal evidence that the Paris Agreement led to a repricing of corporate credit driven primarily by firms' exposure to carbon-intensive production networks. While firms with higher direct emissions appear to face higher borrowing costs when considered in isolation, this effect is not robust once network exposure is taken into account. Instead, both empirical strategies indicate that lenders respond most strongly to the carbon intensity embedded in firms' upstream supply chains, reflecting the transmission of transition risk through production networks rather than through firms' own emissions alone.

The sharp timing of the response around the entry into force of the Paris Agreement, combined with the absence of differential pre-trends, supports a causal interpretation in which climate policy expectations are rapidly internalized by credit markets. Borrowing costs increase disproportionately for firms positioned downstream of carbon-intensive suppliers, even when their own direct emissions remain low. These findings underscore that climate transition risk operates as a systemic and forward-looking phenomenon and that credit markets price not only firm-level environmental footprints, but also inherited exposure to upstream carbon risk along the supply chain.

4 Conclusion and Further Research

This paper has developed a theoretical framework that explains how environmental risks are transmitted into credit markets. Building on the theory of production networks, we constructed a static general equilibrium model with heterogeneous firms, consumers, and competitive banks. Firms are exposed to both physical and transition risks, and must commit to input and labor allocations before shocks realize, which generates the possibility of default. Because production relies on intermediate goods, shocks do not stop at the firm level but propagate through supply chains. The Leontief inverse provides the natural representation of this mechanism: primitive shocks are transformed into network-adjusted shocks that affect all downstream firms. Banks, being residual claimants, set interest rates so as to break even in expectation. This mechanism implies that borrowing costs emerge endogenously as an equilibrium outcome of direct and inherited exposures, independent of productivity differences. The model shows that even firms with low direct emissions face higher interest rates when embedded in brown supply chains, and that borrowing costs for brown firms are systematically higher, consistent with greater transition risk. In particular, the model clarifies how systemic risk, network propagation, and endogenous pricing jointly determine the allocation of credit.

We then validated these predictions empirically using U.S. data. Combining firm-level borrowing costs and balance sheet information from Compustat with sector-level CO_2 emissions from the Environmental Protection Agency and detailed input-output linkages from national accounts, we constructed measures of both direct and network-adjusted carbon exposure at the six-digit NAICS industry level and assigned them to firms operating within those sectors. Using a panel covering the period 2011 to 2023, we document that firms exposed to higher carbon intensity face significantly higher borrowing costs. When considered separately, both direct emissions and network-adjusted emissions are positively associated with firms' cost of debt.

More importantly, once both measures are included jointly, network-adjusted emissions fully account for the observed relationship between carbon exposure and borrowing costs. The effect of direct emissions becomes economically small and statistically insignificant, while exposure to carbon-intensive production networks remains large and highly significant. This finding indicates that lenders primarily price firms' inherited exposure to transition risk transmitted through supply chains, rather than firms' own direct emissions. Even firms operating in relatively clean industries face higher borrowing costs when embedded in carbon-intensive upstream networks, consistent with the model's prediction that transition risk is a systemic, network-amplified phenomenon.

To establish causality, we exploit the 2016 entry into force of the Paris Agreement as a quasi-natural experiment that generated a common and externally imposed shift in climate policy expectations. Difference-in-differences estimates show that borrowing costs increased disproportionately for firms operating in carbon-intensive production networks following the Agreement, while firms' own direct emissions play no independent role once network exposure is taken into account. Event-study results further support this interpretation: borrowing costs for firms in brown networks rise sharply at the time of the policy shock, with no evidence of differential pre-trends, and the post-treatment effects are statistically significant only for network-based exposure. Together, these results provide causal evidence that credit markets internalize climate transition risk primarily through production networks rather than through firms' direct emissions alone.

Together, our contributions are theoretical and empirical. Theoretically, we provide the first general equilibrium model of climate-related credit risk that embeds transition shocks into a production network framework with default. Empirically, we demonstrate that credit markets in the U.S. already internalize environmental risks, pricing not only firm-level emissions but also upstream exposures, and we provide causal evidence that transition risks are amplified by supply chains.

The broader implication is that environmental risk is systemic and not idiosyncratic. Borrowing costs depend as much on the greenness of a firm's network as on its own emissions. This finding supports recent policy initiatives, such as the European Commission's Corporate Sustainability Due Diligence Directive EC (2024), which explicitly accounts for supply chain-related risks. Moreover, our results suggest that market-based incentives can reinforce climate policy: by penalizing brown firms with higher capital costs and rewarding greener firms and supply chains, credit markets can accelerate the reallocation of resources toward a low-carbon economy.

Finally, this research also provides a foundation for the design of policy interventions. In an extension of the model, in Appendix C, we introduce a central bank facility with collateralized loans based on haircuts. We show that optimally chosen haircuts on brown-sector collateral can reduce the variance of financial risk without distorting real allocations by redistributing exposures between private banks and the central bank. This result highlights how prudential policy can be designed to internalize the network-amplified transition risk and stabilize the financial system. By linking production networks, environmental shocks, and monetary intermediation, our framework thus opens the way for future research on the effectiveness of green credit policies and their role in fostering a more resilient and sustainable economy.

References

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

Altavilla, C., M. Boucinha, M. Pagano, and A. Polo (2023): "Climate risk, bank lending and monetary policy," *ECB Working Paper Series*.

Baker, M., D. Bergstresser, G. Serafeim, and J. Wurgler (2018): "Financing the response to climate change: The pricing and ownership of U.S. green bonds," *NBER Working Paper*.

Bartram, Sohnke.M., H. K. and K. Sehoon (2022): "Real effects of climate policy: financial constraints and spillovers," *Journal of Financial Economics*, 143, 668–696.

Becker, R. A., W. B. Gray, and J. Marvakov (2021): "NBER-CES Manufacturing Industry Database (1958-2018, version 2021a)," *National Bureau of Economic Research*.

Bolton, P. and M. Kacperczyk (2021): "Do Investors Care about Carbon Risk?" *Journal of Financial Economics*, 517–549.

Chava, S. (2014): "Environmental Externalities and Cost of Capital," *Management Science*, 2223–2247.

Como, G., F. Fagnani, E. Luciano, A. Milazzo, and M. Scarsini (2026): "Rigidity and default in production networks," *Working Paper available on SSRN*.

Dechezlepretre, A., D. Nachtigall, and F. Venmans (2021): "The Joint Impact of the European Union Emissions Trading System on Carbon Emissions and Economic Performance," *Journal of Environmental Economics and Management*, 110, 102570.

Delis, M. D., K. de Greiff, and S. Ongena (2024): "Being stranded with fossil fuel reserves? Climate policy risk and the pricing of bank loans," *Financial Markets, Institutions and Instruments*.

EC (2024): "CSDD (Corporate Sustainability Due Diligence Directive)," European Commission Directive 2024/1760.

Ehlers, T., F. Packer, and K. D. Greiff (2022): "The Pricing of Carbon Risk in Syndicated Loans: Which Risks Are Priced and Why?" *Journal of Banking & Finance*.

Gormsen, N. J. and C. Huber (2023): "Climate Capitalists," *Working Paper.*

Goss, A. and G. S. Roberts (2011): "The impact of corporate social responsibility on the cost of bank loans," *Journal of Banking and Finance*, 1794–1810.

Ivanov, I., M. Kruttli, and K. L. Watugala (2023): "Banking on Carbon: Lending to Carbon-Intensive Borrowers," *Review of Financial Studies*, 1640–1684.

Kacperczyk, M. and J. L. Peydro' (2024): "Carbon emissions and the bank-lending channel," *ECGI working paper 991/2024.*

Kruse, T., M. Mohnen, and M. Sato (2024): "Do Financial Markets Respond to Green Opportunities?" *Journal of the Association of Environmental and Resource Economists.*

Pastor, L., R. F. Stambaugh, and L. A. Taylor (2021): "Sustainable Investing in Equilibrium," *Journal of Financial Economics*, 550–571.

Pellet, T. and A. Tahbaz-Salehi (2023): "Rigid Production Networks," *Journal of Monetary Economics*, 86–102.

Pittman, J. A. and S. Fortin (2003): "Auditor choice and the cost of debt capital for newly public firms," *Journal of Accounting and Economics.*

Porter, M. E. (1991): "Americaâs Green Strategy," *Scientific American*, 264, 168–173.

Porter, M. E. and C. van der Linde (1995): "Toward a New Conception of the Environment–Competitiveness Relationship," *Journal of Economic Perspectives*, 9, 97–118.

Reghezza, A., Y. Altunbas, D. Marques-Ibanez, A. Maddaloni, and M. Spaggiari (2022): "The Impact of Climate Change on the European Financial System: Credit and Market Channels," *ECB Working Paper Series.*

Salman, M., X. Long, G. Wang, and D. Zha (2022): "Paris climate agreement and global environmental efficiency: New evidence from fuzzy regression discontinuity design," *Energy Policy.*

Solow, R. M. (1957): "Technical Change and the Aggregate Production Function," *The Review of Economics and Statistics.*

Zerbib, O. D. (2019): "The effect of pro-environmental preferences on bond prices: Evidence from green bonds," *Journal of Banking and Finance*, 39–60.

Appendix

A Proofs

Proof of theorem 1

Proof. Recall that primitive shocks η to green firms dominate those to brown firms in first-order stochastic dominance. For brevity, write $\eta_g \geq \eta_b$. This implies:

$$P(\eta_g \leq x) \leq P(\eta_b \leq x) \quad \text{for all } x \tag{47}$$

where P is the probability on Ω. Since $\rho = L\eta$, the same dominance holds for total (network-propagated) shocks:

$$P(\rho_g \leq Lx) \leq P(\rho_b \leq Lx). \tag{48}$$

Applying monotonic transformations:

$$P\left(\exp(\rho_g) \leq \exp(Lx)\right) \leq P\left(\exp(\rho_b) \leq \exp(Lx)\right). \tag{49}$$

Dividing by $\mathbb{E}[\exp(\rho_b)]$ preserves the inequality:

$$P\left(\frac{\exp(\rho_g)}{\mathbb{E}[\exp(\rho_b)]} \leq \frac{\exp(Lx)}{\mathbb{E}[\exp(\rho_b)]}\right) \leq P\left(\frac{\exp(\rho_b)}{\mathbb{E}[\exp(\rho_b)]} \leq \frac{\exp(Lx)}{\mathbb{E}[\exp(\rho_b)]}\right). \tag{50}$$

Using the fact that first-order stochastic dominance implies $\mathbb{E}[u(\rho_g)] \geq \mathbb{E}[u(\rho_b)]$ for all non-decreasing functions u, and taking $u = \exp$, we also have:

$$P\left(\frac{\exp(\rho_g)}{\mathbb{E}[\exp(\rho_g)]} \leq \frac{\exp(Lx)}{\mathbb{E}[\exp(\rho_b)]}\right) \leq P\left(\frac{\exp(\rho_g)}{\mathbb{E}[\exp(\rho_b)]} \leq \frac{\exp(Lx)}{\mathbb{E}[\exp(\rho_b)]}\right) \tag{51}$$

and therefore

$$P\left(\frac{\exp(\rho_g)}{\mathbb{E}[\exp(\rho_g)]} \leq \frac{\exp(Lx)}{\mathbb{E}[\exp(\rho_b)]}\right) \leq P\left(\frac{\exp(\rho_b)}{\mathbb{E}[\exp(\rho_b)]} \leq \frac{\exp(Lx)}{\mathbb{E}[\exp(\rho_b)]}\right). \tag{52}$$

for all values of x and therefore all non negative real values that obtain from $\frac{\exp(Lx)}{\mathbb{E}[\exp(\rho_b)]}$.

This shows that normalized shocks to green firms dominate those to brown firms:

$$\tau_g \geq \tau_b. \tag{53}$$

If the conditions in the theorem are true, then:

$$\epsilon_g \geq \epsilon_b. \tag{54}$$

Since the min is weakly increasing in both arguments, it follows that:

$$\mathbb{E}[\min(\tau_g, \epsilon_g)] \geq \mathbb{E}[\min(\tau_b, \epsilon_b)], \tag{55}$$

which implies:

$$r_g = \frac{1}{\mathbb{E}[\min(\tau_g, \epsilon_g)]} - 1 \leq \frac{1}{\mathbb{E}[\min(\tau_b, \epsilon_b)]} - 1 = r_b. \tag{56}$$

for any brown firm $b > n_1$ and green firm $g \leq n_1$.

■

Proof of Theorem 2

Proof. We established in the previous theorem that the first order stochastic dominance of the primitive shocks is preserved by the total shocks

$$\tau_g \geq \tau_b. \tag{57}$$

Call m the minimum of the support of the total shocks to green sectors, and M the same quantity for brown sectors,

$$M = \min_{j \leq n_1} \mathbf{support}(\tau_{\mathbf{j}}) \tag{58}$$

$$m = \min_{j > n_1} \mathbf{support}(\tau_{\mathbf{j}}) \tag{59}$$

and observe that

$$\sum_{j \le n_1} A_{gj} \tau_j \ge M \sum_{j \le n_1} A_{gj} \tag{60}$$

$$\sum_{j > n_1} A_{gj} \tau_j \ge m \sum_{j > n_1} A_{gj} \tag{61}$$

$$m \le M \tag{62}$$

$$\epsilon_g \ge M \sum_{j \le n_1} A_{gj} + m \sum_{j > n_1} A_{gj} + \beta_g \tag{63}$$

$$\epsilon_b \ge M \sum_{j \le n_1} A_{bj} + m \sum_{j > n_1} A_{bj} + \beta_b \tag{64}$$

If $\sum_{j \le n_1} A_{gj} \ge \sum_{j \le n_1} A_{bj}$ and $\sum_{j > n_1} A_{gj} = \sum_{j > n_1} A_{bj}$, $\beta_g \ge \beta_b$, then the lower bound for ϵ_g is greater than the one for $\ge \epsilon_b$ and therefore $min(\epsilon_g, \tau_g) \ge \min(\epsilon_b, \tau_b)$. The expectation follows the same inequality, and therefore $r_g \le r_b$.

■

Proof of Corollary 1

Proof. Under the assumptions of the lemma, $\sum_{j \le n_1} A_{bj} = 0$ and $\sum_{j > n_1} A_{gj} = 0$ for all indices $g \le n_1, b > n_1$,. The conclusion is straightforward. ■

Proof of Corollary 2

Proof. Under the stated assumptions, $\tau_g = 1, \epsilon_g = 1$. The conclusion is straightforward. ■

B Empirical Analysis

B.1 CO_2 Emissions Data

This section of the appendix provides a detailed and rigorous description of the CO_2 emissions data sources, transformations, and calibration procedures used in the paper. The objective is to map the theoretical objects defined in the model to empirically observable counterparts in a manner that is transparent, internally consistent, and replicable. We describe in turn the construction of direct and network-adjusted emissions and the calibration of the production network matrix used in the model.

Throughout the appendix, industries are indexed by $k \in \mathcal{V}$, consistently with the notation adopted in the theoretical and empirical sections.

B.1.1 EPA Greenhouse Gas Reporting Program

The primary source of direct emissions data is the U.S. Environmental Protection Agency (EPA) Greenhouse Gas Reporting Program (GHGRP). The GHGRP was established in 2008 under the authority of the Clean Air Act with the objective of providing comprehensive, facility-level information on greenhouse gas emissions across the U.S. economy. The program is designed to support federal and state climate policy, improve emissions transparency, and provide a consistent statistical foundation for the analysis of greenhouse gas emissions across industries and over time.

Under the GHGRP, reporting is mandatory for facilities operating in the United States whose annual direct greenhouse gas emissions exceed 25,000 metric tons of carbon dioxide equivalent (CO_2e). In addition to large stationary sources, the program also covers selected upstream suppliers of fossil fuels and industrial gases whose products give rise to downstream emissions. As a result, the GHGRP captures the vast majority of emissions generated by energy-intensive activities, including manufacturing, electricity generation, petroleum refining, chemical production, metals, minerals, and other heavy industrial sectors. While smaller facilities fall below the reporting threshold, the GHGRP is widely recognized as covering a dominant share of industrial greenhouse gas emissions in the United States.

Emissions are reported annually at the facility level and encompass all major greenhouse gases regulated under the program, including carbon dioxide (CO_2), methane (CH_4), nitrous oxide (N_2O), and fluorinated gases. Reported quantities of individual gases are converted into a common metric of carbon dioxide equivalent using 100-year global warming potentials consistent with the Intergovernmental Panel on Climate Change Fourth Assessment Report. This conversion allows emissions from different gases to be aggregated into a single, economically meaningful measure of climate impact. Emissions are expressed in metric tons of CO_2 equivalent and reflect direct emissions generated by on-site fuel combustion, industrial processes, and other covered activities.

A key feature of the GHGRP is the detailed industry classification reported for each facility. Each reporting entity is assigned a Primary NAICS Code, corresponding to the six-digit North American Industry Classification System. This classification reflects the facility's main economic activity and provides a natural bridge between facility-level emissions data and industry-level economic data. The availability of consistent NAICS identifiers makes it possible to aggregate emissions across facilities to construct sector-level emissions measures that align with standard input–output and national accounts data.

In this paper, we use the EPA-provided multi-year summary file `ghgp_data_by_year_2023.xlsx`, which consolidates facility-level GHGRP data for reporting years 2010 through 2023. Within this file, we rely on the worksheet labeled Direct Point Emitters, which contains annual total direct emissions for each reporting facility along with its Primary NAICS Code and other identifying information. For each facility and year, the dataset reports the total quantity of direct greenhouse gas emissions, expressed in metric tons of CO_2 equivalent.

The GHGRP is particularly well suited for the analysis conducted in this paper for several reasons. First, it provides emissions data at a high level of disaggregation, allowing emissions to be mapped precisely to six-digit NAICS industries. Second, the use of a consistent reporting threshold and standardized measurement methodology ensures comparability across facilities and over time. Third, the focus on direct emissions aligns closely with the concept of sector-specific productivity shocks in the model, which operate through industries' own production activities. Finally, the annual frequency of the data allows emissions measures to be matched directly to the timing of financial variables and policy events considered in the empirical analysis.

Aggregation to Six-Digit NAICS Industries Facility-level emissions are aggregated to the six-digit NAICS industry level in order to align emissions data with the production network used in the model. This choice is primarily dictated by the structure of the BEA input–output tables, which represent inter-industry linkages at the level of detailed industries that can be consistently mapped to six-digit NAICS classifications. As a result, the six-digit

NAICS level constitutes the most granular level at which direct emissions from the EPA GHGRP can be merged with the BEA production network in a consistent and economically meaningful manner.

Let $e_{f,t}$ denote the total direct greenhouse gas emissions of facility f in year t, measured in metric tons of CO_2 equivalent. Each facility is assigned to an industry $k(f)$ based on its Primary NAICS Code reported in the GHGRP. Aggregating emissions across facilities within the same six-digit NAICS industry yields a sector-level measure of direct emissions.

Formally, sector-level direct emissions are constructed as

$$\texttt{CO2}_{k,t} = \sum_{f \in k} e_{f,t}, \tag{65}$$

where the summation runs over all facilities whose primary NAICS code corresponds to sector k. This aggregation preserves the full coverage of facility-level emissions while producing industry-level measures that are directly comparable across sectors and over time.

The aggregation is implemented in the Python script `0_CO2_Extraction.py`. The script programmatically identifies the header row of the EPA spreadsheet, extracts all year-specific emissions columns, reshapes the data to long format, and aggregates emissions across facilities within each six-digit NAICS sector and year. The resulting panel dataset contains one observation per industry k and year t, with $\texttt{CO2}_{k,t}$ expressed in metric tons of CO_2e.

Aggregating emissions at the six-digit NAICS level ensures consistency with the BEA input–output tables used to construct the production network and allows emissions to be propagated through the economy using the Leontief framework described in the subsequent subsection. In the empirical analysis, emissions are rescaled and reported in million metric tons of CO_2e for ease of interpretation, without affecting the underlying construction.

B.1.2 BEA Input-Output Tables

To capture production network linkages, we rely on industry-by-industry input–output tables produced by the U.S. Bureau of Economic Analysis. Specifically, we use the total requirements matrices provided in the file `IxI_TR_Detail.xlsx`, which contains benchmark-year tables for 2007, 2012, and 2017.

For a given benchmark year y, the BEA total requirements matrix L_y is defined such that each element $L_{jk,y}$ measures the total dollar value of output from industry j required, directly and indirectly, to deliver one dollar of final demand for the output of industry k. These matrices therefore capture the full upstream propagation of production through the economy.

The Python script `0_IO_Tables_Extraction.py` extracts the numerical blocks corresponding to the total requirements matrices and the associated industry codes. Rows and columns corresponding to subtotals or non-industry entries are removed. The resulting matrices are square and indexed by industry codes that can be mapped to six-digit NAICS sectors.

B.1.3 From BEA Total Requirements to Model Production Networks

This subsection explains in detail how the production network matrix used in the model is constructed from the BEA industry-by-industry total requirements tables and clarifies the distinction between empirical input–output objects and their model-consistent counterparts.

The BEA industry-by-industry total requirements table for a given benchmark year y is a square matrix, denoted $\mathbf{L}_y^{TR}$. In the Excel file `IxI_TR_Detail.xlsx`, each table appears as a rectangular numerical block in which rows correspond to supplying industries and columns correspond to industries delivering final demand. The Python script `0_IO_Tables_Extraction.py` reads the industry labels from the first row and first column of this block and constructs a square matrix indexed consistently across rows and columns. An entry $L_{jk,y}^{TR}$ therefore corresponds to industry j as a supplier (row) and industry k as a demander (column).

Each column k of $\mathbf{L}_y^{TR}$ reports the total gross output required from all industries in order to deliver one dollar of final demand for the output of industry k. In particular, the element $L_{jk,y}^{TR}$ measures the total dollar value of output from industry j that is required, both directly and indirectly through all upstream production stages, to produce one dollar of final demand for industry k. This column-wise interpretation is crucial for mapping the BEA object into the production network structure of the model.

In the model, the importance coefficient A_{jk} captures the share of sector j's output used as an intermediate input in the production of sector k. Columns of the matrix $\mathbf{A}$ therefore describe the input composition of sector k, while rows describe how the output of a given sector j is distributed across downstream users. Under the standard Leontief production framework, the relationship between the BEA total requirements matrix and the matrix of empirical direct input coefficients can be written as

$$\mathbf{L}_y^{TR} = \left(\mathbf{I} - \mathbf{A}_y^{TR}\right)^{-1}, \tag{66}$$

where $\mathbf{A}_y^{TR}$ denotes the matrix of direct requirements implied by the BEA data and $\mathbf{I}$ denotes the identity matrix. This relationship reflects the fact that total requirements incorporate not only direct input needs but also all higher-order indirect requirements arising from upstream production linkages.

Given the BEA total requirements matrix, the implied empirical direct requirements matrix is therefore recovered as

$$\mathbf{A}_y^{TR} = \mathbf{I} - \left(\mathbf{L}_y^{TR}\right)^{-1}. \tag{67}$$

This inversion step maps the observed BEA production network into a matrix of empirical input shares. In the implementation, the script `0_IO_CO2_Merge.py` loads the cleaned numerical block for $\mathbf{L}_y^{TR}$, verifies that row and column labels coincide, computes its numerical inverse, and constructs $\mathbf{A}_y^{TR}$ accordingly.

The matrix $\mathbf{A}_y^{TR}$ does not yet correspond to the production network used in the model, as the latter incorporates financial frictions that distort observed input shares. Following the calibration procedure described in Appendix B.1.6, we obtain the model-consistent production network matrix $\mathbf{A}_y$ by rescaling $\mathbf{A}_y^{TR}$ using sector-specific interest rate derived from observed borrowing costs.

A key theoretical restriction in the model, stemming from the constant returns-to-scale normalization of industries' production functions, is that the total value of intermediate inputs used by any industry cannot exceed the value of its gross output. In matrix terms, this restriction implies that no column sum of the matrix $\mathbf{A}_y$ exceeds one. As a consequence, $\mathbf{A}_y$ is a column-substochastic matrix, and its spectral radius, defined as the largest modulus of its eigenvalues, is weakly less than one. Throughout the analysis, we assume that the spectral radius of $\mathbf{A}_y$ is strictly less than one. This is a mild and economically natural condition: it is satisfied, for example, if every industry is connected, possibly through multiple production stages, to at least one sector that employs a primary factor such as labor.

Under this assumption, the economy's Leontief inverse used in the model is defined as

$$\mathbf{L}_y = \left(\mathbf{I} - \mathbf{A}_y'\right)^{-1}, \tag{68}$$

where $\mathbf{A}_y'$ denotes the transpose of $\mathbf{A}_y$. The transpose appears because, in the model, columns of $\mathbf{A}_y$ index users while rows index suppliers, so that the propagation of shocks through upstream linkages naturally involves $\mathbf{A}_y'$. The matrix $\mathbf{L}_y$ can be expressed as the convergent Neumann series

$$\mathbf{L}_y = \sum_{h=0}^{+\infty} \left(\mathbf{A}_y'\right)^h. \tag{69}$$

This representation makes explicit that the Leontief inverse aggregates the effects of all production paths of arbitrary length in the network. Each term $\left(\mathbf{A}_y'\right)^h$ captures the contribution of paths of length h, and the nonnegativity of $\mathbf{A}_y$ implies that all entries of

$\mathbf{L}_y$ are nonnegative. In particular, $L_{kj,y}$ measures the cumulative importance of industry j as an upstream supplier to industry k through all possible chains of intermediate input relationships.

The Python script `0_IO_CO2_Merge.py` implements these transformations directly. Starting from the BEA total requirements matrix $\mathbf{L}_y^{TR}$, the script computes $\mathbf{A}_y^{TR}$, rescales it to obtain the model-consistent matrix $\mathbf{A}_y$, and verifies that the resulting matrix satisfies the column-substochasticity and nonnegativity conditions required for the existence of the Leontief inverse. This construction ensures that the empirical production network used in the paper is fully consistent with the theoretical structure of the model and allows productivity and emissions shocks to propagate through the economy according to the Leontief expansion described in Section 2.

B.1.4 Visualization of the Production Network

Figures 3–6 provide graphical representations of the model-consistent production network constructed following the procedure described above. These figures are intended as descriptive illustrations of the structure and intensity of production linkages among manufacturing industries and are not used directly in the quantitative analysis.

Figures 3 and 4 display heatmap representations of the Leontief inverse matrices $\mathbf{L}_{2012}$ and $\mathbf{L}_{2017}$ for manufacturing industries. Each cell (k, j) is colored according to the magnitude of the element $L_{kj,y}$, which measures the cumulative upstream importance of industry j for the production of industry k, aggregating direct and indirect input requirements along all production paths in the network. Brighter colors correspond to stronger total requirements. The prominent diagonal reflects the fact that own-industry effects dominate total requirements, while off-diagonal clusters highlight economically meaningful upstream-downstream linkages across manufacturing sectors. Differences in color intensity patterns across years reflect changes in the structure and strength of production interdependencies over time.

Figures 5 and 6 provide network visualizations of the corresponding model-consistent input-output matrices $\mathbf{A}_{2012}$ and $\mathbf{A}_{2017}$. Nodes represent manufacturing industries, and directed edges capture intermediate input linkages from supplying to using industries. Node size and color encode eigenvector centrality, which measures an industry's importance in the production network by accounting for both the number and the strength of its connections, as well as the centrality of its neighbors. Lighter and larger nodes correspond to industries that are more central in the propagation of shocks through the network. Edge thickness is proportional to the magnitude of the underlying input coefficient, while the absence of labels emphasizes the global structure of the network rather than individual sector identities.

Taken together, these figures illustrate two complementary perspectives on the production network. The Leontief heatmaps emphasize the cumulative strength of upstream dependencies embedded in the economy, while the network diagrams highlight the heterogeneous centrality of industries and the sparsity of economically relevant linkages. Both representations are fully consistent with the theoretical structure of the model and with the construction of the production network matrix described in this appendix.

Figure 3: Leontief Inverse Matrix, Manufacturing Industries, 2012

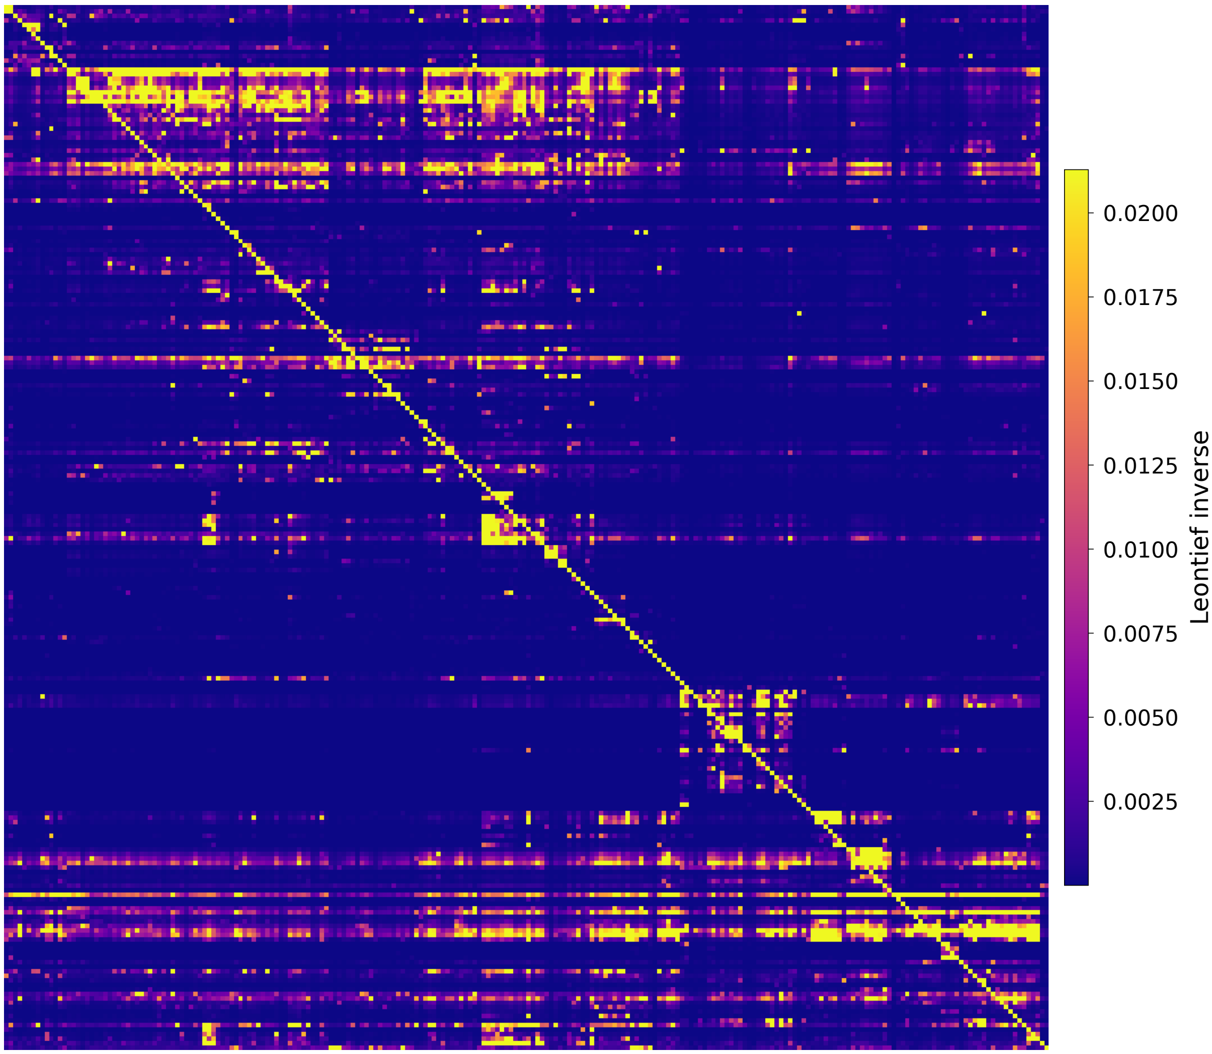

Notes: This figure shows a heatmap of the model-consistent Leontief inverse $\mathbf{L}_{2012} = (\mathbf{I} - \mathbf{A}'_{2012})^{-1}$ restricted to manufacturing industries. Each cell (k, j) is colored according to the magnitude of $L_{kj,2012}$, which measures the cumulative importance of industry j as an upstream supplier to industry k through all direct and indirect input linkages. Brighter colors indicate stronger total (direct and indirect) input requirements.

Figure 4: Leontief Inverse Matrix, Manufacturing Industries, 2017

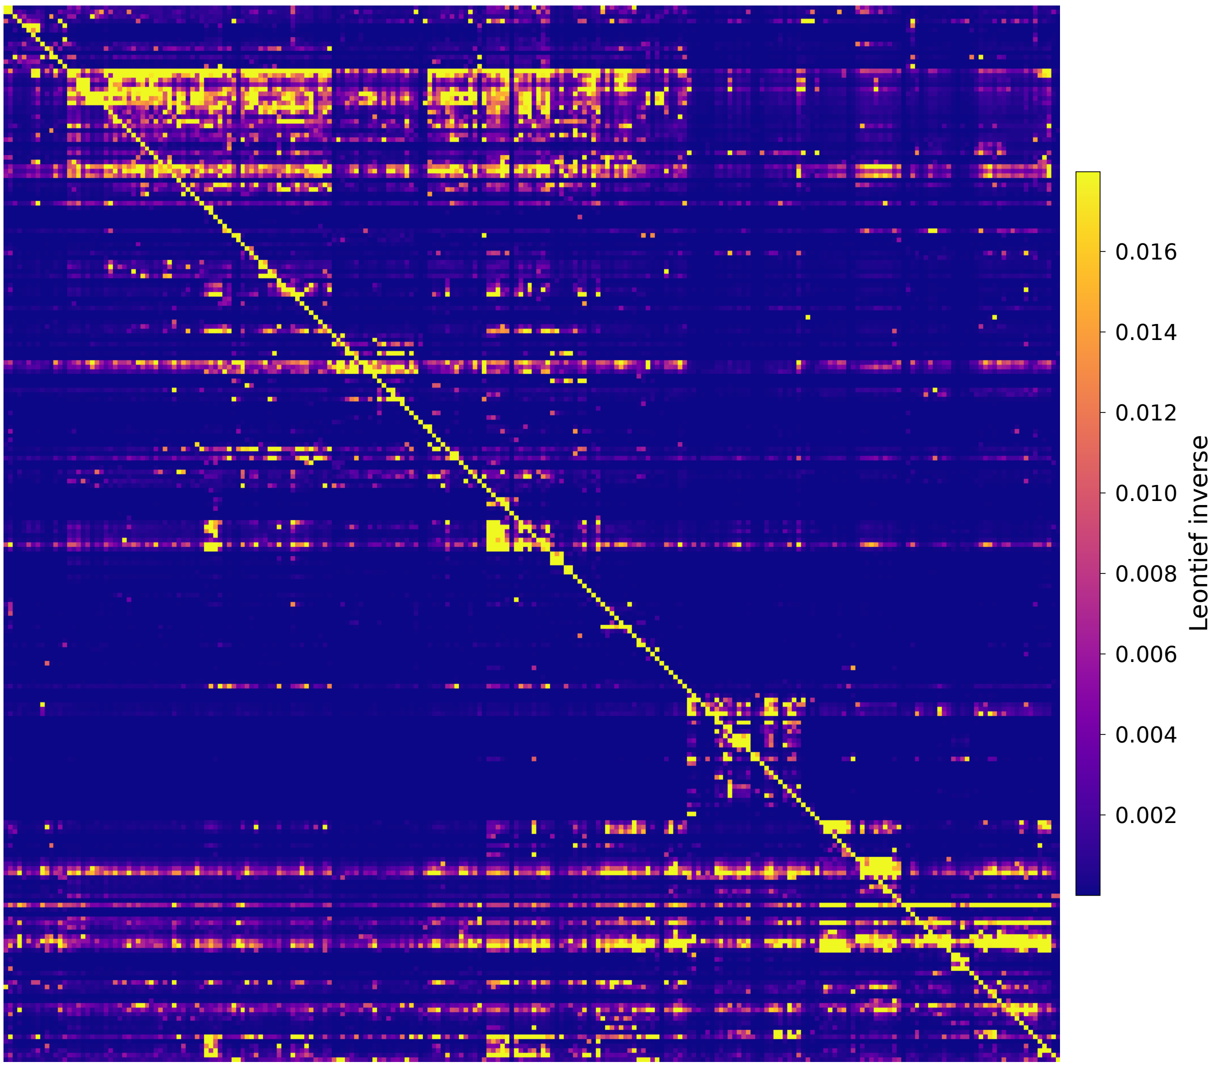

Notes: This figure shows a heatmap of the model-consistent Leontief inverse $\mathbf{L}_{2017} = (\mathbf{I} - \mathbf{A}'_{2017})^{-1}$ restricted to manufacturing industries. Cell colors reflect the magnitude of $L_{kj,2017}$, with brighter cells corresponding to stronger cumulative upstream dependence of industry k on industry j. Differences in the pattern and intensity of off-diagonal elements relative to Figure 3 illustrate changes in the structure of manufacturing production linkages between 2012 and 2017.

Figure 5: Production Network Representation, Manufacturing Industries, 2012

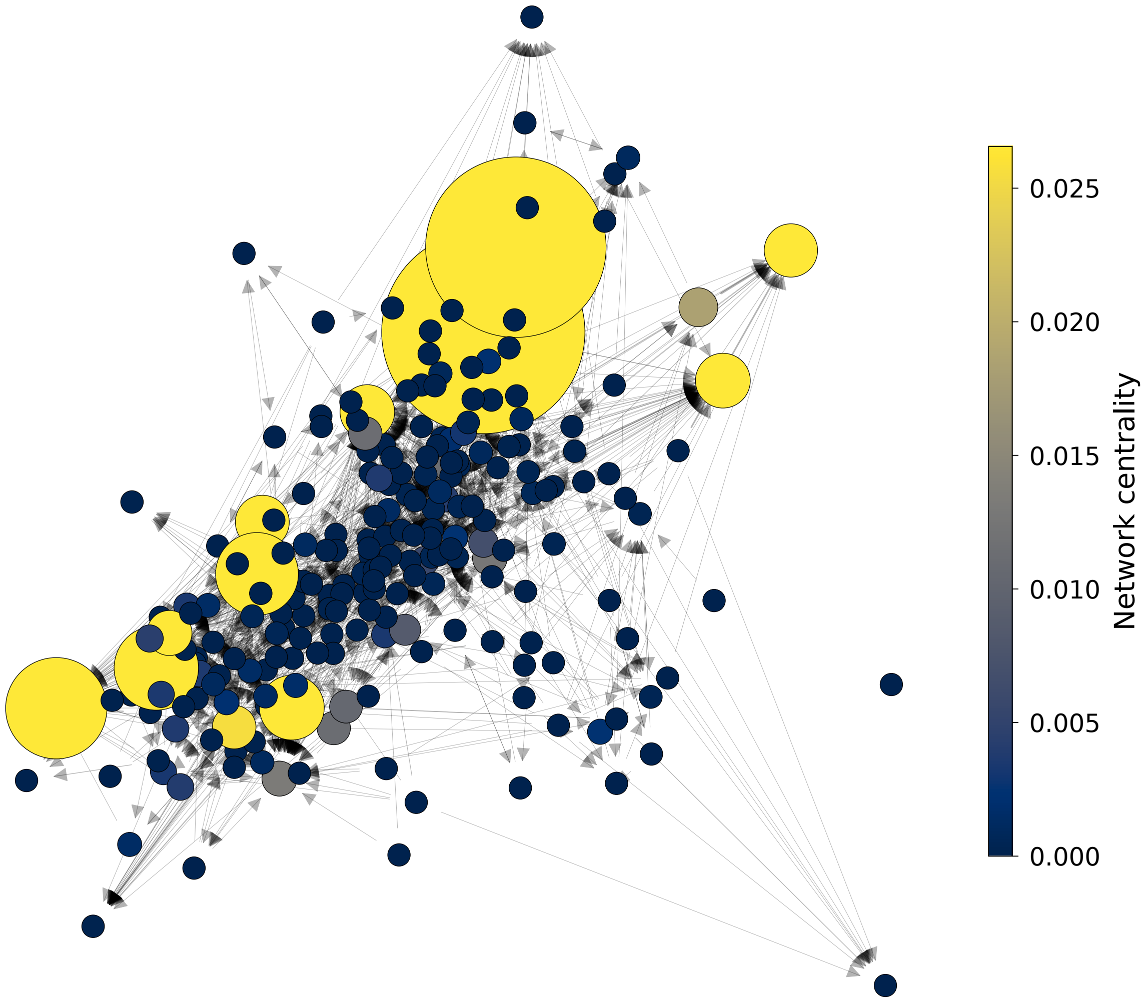

Notes: Nodes represent manufacturing industries and directed edges represent intermediate input linkages in the model-consistent production network $\mathbf{A}_{2012}$. Node size and color encode eigenvector centrality, so that larger and lighter nodes are more central in the network, taking into account both the number and the strength of their connections and the centrality of their neighbors. Edge thickness is proportional to the corresponding input coefficient. The absence of labels emphasizes the overall topology of the production network rather than individual sector identities.

Figure 6: Production Network Representation, Manufacturing Industries, 2017

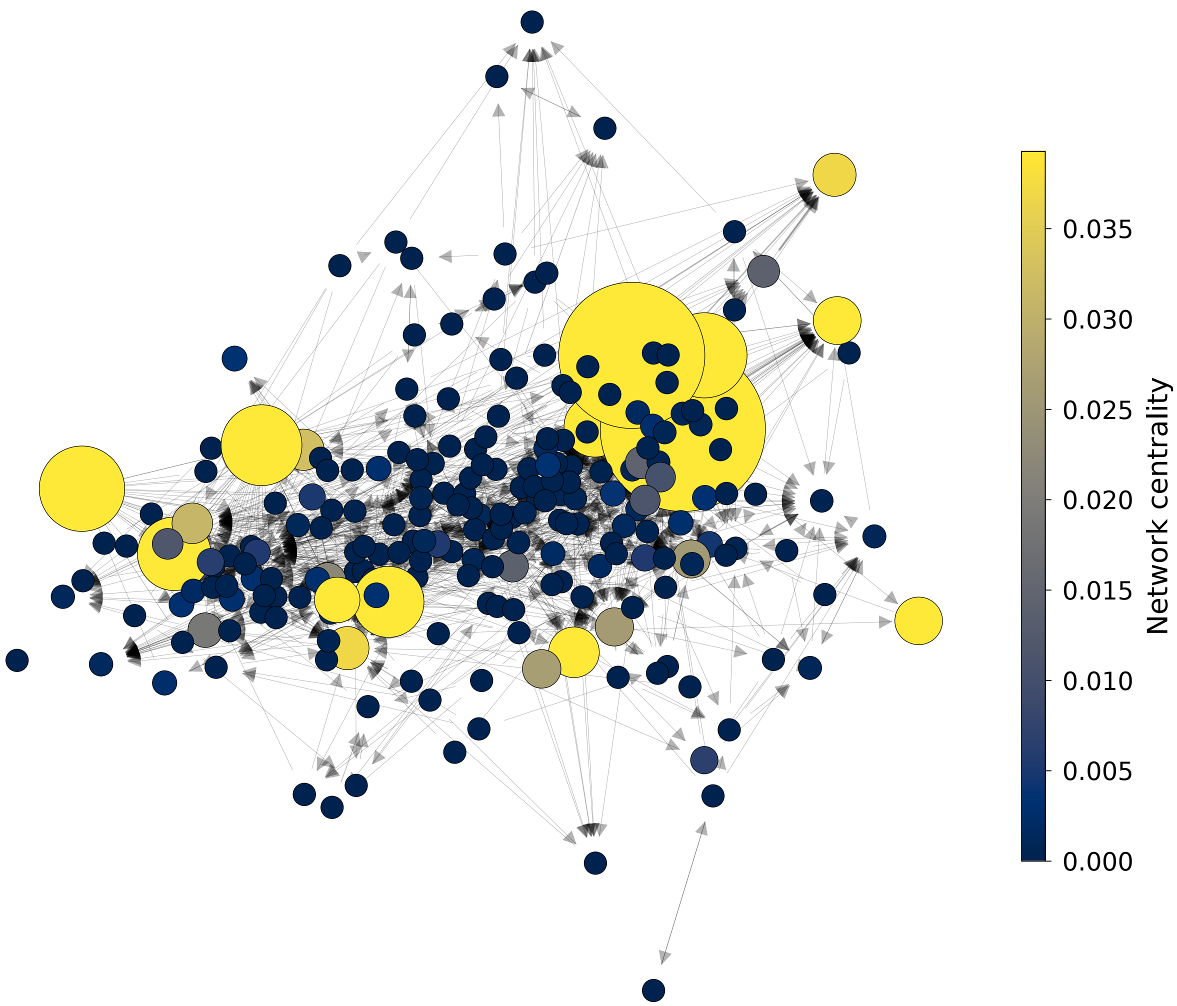

Notes: This figure provides the network representation of the model-consistent production matrix $\mathbf{A}_{2017}$ for manufacturing industries. As in Figure 5, node size and color reflect eigenvector centrality and edge thickness reflects the strength of intermediate input linkages. Comparing Figures 5 and 6 highlights changes in the centrality of industries and in the density and orientation of production linkages over time.

B.1.5 Construction of the Network-Adjusted Emissions Matrix

This subsection describes how the calibrated production network matrix is combined with direct emissions data to construct the network-adjusted emissions variable `Network CO2`. The objective is to capture the total carbon exposure embodied in each industry's production, accounting not only for emissions generated directly by the industry itself, but also for emissions generated by all upstream suppliers throughout the production network.

Let $\texttt{CO2}_t$ denote the vector of direct greenhouse gas emissions across industries in year t. Each element $\texttt{CO2}_{k,t}$ corresponds to the total direct emissions of industry k, measured in metric tons of CO_2 equivalent, and is constructed from the EPA GHGRP as described in Appendix B.1.1. In the empirical dataset, this object is stored as the variable `CO2`. The vector $\texttt{CO2}_t$ is ordered consistently with the rows and columns of the model-consistent production network matrix $\mathbf{A}_t$ and its associated Leontief inverse $\mathbf{L}_t$.

We construct network-adjusted emissions using a transformation, treating direct emissions as a primitive source of environmental exposure that propagates through the production network.

Formally, the vector of network-adjusted emissions in year t, stored in the empirical dataset as the variable `Network CO2`, is defined as $\texttt{Network CO2}_t = \mathbf{L}_t' \texttt{CO2}_t$ where $\mathbf{L}_t'$ denotes the transpose of the model Leontief inverse corresponding to the benchmark system used for year t. The transpose appears for the same reason as in the theoretical analysis: columns of the production network matrix index input users, while rows index suppliers, so upstream exposure is captured by premultiplication with $\mathbf{L}_t'$.

Element by element, the network-adjusted emissions of industry k in year t are given by

$$\texttt{Network CO2}_{k,t} = \sum_{j \in \mathcal{V}} L_{jk,t} \, \texttt{CO2}_{j,t}. \tag{70}$$

This expression makes explicit that $\texttt{Network CO2}_{k,t}$ aggregates the direct emissions of all upstream industries j, weighted by their cumulative importance in the production of sector k through all direct and indirect input linkages. The Leontief inverse therefore accounts for emissions embodied in intermediate inputs of arbitrary order, including emissions generated by suppliers of suppliers and so on.

In practice, BEA industry-by-industry total requirements tables are only available for a limited set of benchmark years (2007, 2012, 2017). The script `0_IO_CO2_Merge.py` uses the 2012 and 2017 benchmark systems to approximate the production network over our sample period. Specifically, for each calendar year t in the GHGRP data, the script proceeds as follows:

1. It constructs the manufacturing Leontief inverse for 2012 and 2017, denoted $\mathbf{L}_{2012}^{M}$ and $\mathbf{L}_{2017}^{M}$, from the corresponding calibrated manufacturing direct requirements matrices as described in Appendix B.1.3.

2. For years $t \leq 2014$, it sets $\mathbf{L}_t = \mathbf{L}_{2012}^{M}$ and uses this matrix to construct $\texttt{Network CO2}_t$.

3. For years $t \geq 2015$, it sets $\mathbf{L}_t = \mathbf{L}_{2017}^{M}$ and uses this matrix to construct $\texttt{Network CO2}_t$.

In the code, this logic is implemented by the rule

$$\mathbf{L}_t = \begin{cases} \mathbf{L}_{2012}^{M} & \text{if } t \leq 2014, \\ \mathbf{L}_{2017}^{M} & \text{if } t \geq 2015, \end{cases} \tag{71}$$

and by calling a helper function that, for each year t, aligns the `CO2` data with the index of the chosen Leontief matrix and applies the transformation $\mathbf{L}_t'\texttt{CO2}_t$ only to industries with non-missing emissions.

Because the entries of $\mathbf{L}_t$ are unitless and $\texttt{CO2}_t$ is measured in metric tons of CO_2 equivalent, the resulting vector $\texttt{Network CO2}_t$ is expressed in the same units. In the empirical analysis, both `CO2` and `Network CO2` are sometimes rescaled and reported in million metric tons of CO_2 equivalent for ease of interpretation, without affecting the underlying construction.

This adjusted emissions network constitutes the empirical counterpart to the propagation of shocks described in the theoretical model and provides the basis for distinguishing between direct environmental exposure (captured by `CO2`) and exposure arising through production network linkages (captured by `Network CO2`) in the empirical analysis.

B.1.6 Calibration of the Production Network Matrix

This subsection describes how the empirical input–output matrix is mapped into the production network matrix used in the theoretical model.

In the frictionless benchmark of the model, the first-order conditions imply that the importance coefficient A_{jk} equals the share of inputs from sector j in total output of sector k. In particular, under perfect information and zero interest rates, the model yields

$$A_{jk} = \frac{p_j z_{jk}}{p_k y_k}, \tag{72}$$

where $p_j z_{jk}$ denotes the value of inputs from sector j used by sector k, and $p_k y_k$ denotes the value of total output of sector k. This expression corresponds exactly to the direct requirements coefficients computed from the BEA input–output tables and provides the standard calibration approach adopted in the literature.

In the presence of financial rigidities, however, the model implies that observed input shares are distorted by sector-specific discount factors. Using the first-order conditions of the model, the importance coefficients satisfy

$$\frac{p_j z_{jk}}{p_k y_k} = \frac{A_{jk}}{1 + r_k \theta_k} = A_{jk} \exp(-\zeta_k), \tag{73}$$

where r_k denotes the borrowing rate faced by sector k, θ_k captures the leverage, and $\zeta_k = \log(1 + r_k \theta_k)$. Rearranging yields the calibration equation

$$A_{jk} = \frac{p_j z_{jk}}{p_k y_k} \exp(\zeta_k). \tag{74}$$

Accordingly, the production network matrix used in the model is obtained by rescaling the empirical BEA direct requirements matrix by sector-specific discount factors derived from observed borrowing costs. This adjustment ensures consistency between the theoretical importance coefficients and the empirically observed production structure in the presence of leverage.

B.2 Compustat Firm-Level Financial Data

Firm-level financial variables are obtained from the Compustat Fundamentals Annual database, covering publicly listed U.S. firms. The sample spans the years 2012 to 2023. Firms are included in the analysis only if they report valid six-digit NAICS industry codes and have non-missing entries for interest expense and debt-related variables. Observations with implausible or extreme financial ratios are excluded as described below.

The key financial variables used in our analysis include:

- `xint`: Interest expense
- `dltt`: Long-term debt obligations
- `dlc`: Short-term debt (due within one year)
- `lt`: Total liabilities
- `at`: Total assets
- `che`: Cash and equivalents
- `ebit`: Earnings before interest and taxes
- `ppent`: Net property, plant, and equipment
- `ceq`, `mkvalt`: Variables used to compute market-to-book ratios

To mitigate the influence of outliers, observations with implausible or extreme values of leverage, cash holdings, total assets, and investment rates are trimmed at the top 1 percent of their respective distributions. All financial variables entering the empirical analysis are expressed in logarithms. Observations with missing or non-positive values for key balance-sheet items are excluded. The resulting sample provides a consistent panel of U.S. firms suitable for analyzing firm-level financial conditions using annual Compustat data.

B.2.1 Measuring the Cost of Debt

A central variable in our empirical analysis is the firm-level cost of debt. We construct this measure using firms' realized interest expenses and outstanding debt balances, with the goal of capturing the effective borrowing rate paid on all interest-bearing obligations. This approach reflects the average cost of debt financing faced by a firm in a given year, aggregating across maturities and debt instruments.

Our preferred measure of the cost of debt is defined as the ratio of interest expense to total interest-bearing debt, expressed as a percentage. Formally, for firm i in year t, we define:

$$r_{i,t} = \frac{\texttt{xint}_{i,t}}{\texttt{dltt}_{i,t} + \texttt{dlc}_{i,t}} \times 100, \tag{75}$$

where $\texttt{xint}_{i,t}$ denotes interest expense and $\texttt{dltt}_{i,t} + \texttt{dlc}_{i,t}$ corresponds to total interest-bearing debt, composed of long-term debt due beyond one year and short-term debt due within one year.

This firm-level measure captures the effective interest rate paid on outstanding debt and is well suited to environments in which borrowing costs may vary systematically across firms due to differences in risk exposure, capital structure, or regulatory conditions. Unlike contractual interest rates on individual instruments, this measure reflects the realized financial burden of debt financing faced by the firm as a whole.

Because the construction of implied interest rates using accounting data can generate extreme values due to reporting errors, temporary debt reductions, or timing mismatches between interest expenses and balance sheet stocks, we implement a multi-step data cleaning procedure. First, we drop observations with missing values for interest expense or interest-bearing debt, as well as firms without valid industry identifiers. Second, we winsorize key balance sheet components, including interest expense and debt variables, at the 2.5th and 97.5th percentiles to reduce the influence of extreme accounting realizations before computing the implied rate.

After constructing the implied interest rate, we restrict the sample to observations with values less than or equal to 30 percent. This upper bound eliminates implausibly high rates that are typically driven by very small denominators or transitory accounting irregularities rather than genuine borrowing costs. Following this truncation step, we winsorize the interest rate distribution at the 5th and 95th percentiles, consistent with standard practice in the corporate finance literature, including Pittman and Fortin (2003). This symmetric winsorization further limits the influence of residual outliers while preserving economically meaningful variation in borrowing costs.

The same trimming and winsorization strategy is applied to an alternative interest rate measure based on total liabilities, which serves as a robustness check in the empirical analysis. Together, these filtering steps ensure that the resulting cost of debt measure reflects economically plausible financing conditions and is not driven by mechanical accounting artifacts or extreme observations.

Alternative Measure of Cost of Debt. To assess the robustness of our results, we also construct an alternative measure of the cost of debt based on total liabilities:

$$\mathrm{r}_{i,t}^{\mathrm{Alt}} = \frac{\texttt{xint}_{i,t}}{\texttt{lt}_{i,t}} \times 100. \tag{76}$$

Although this specification may understate the cost of debt by including non-interest-bearing liabilities in the denominator, it is less sensitive to misclassification or missing components of interest-bearing debt and provides a conservative benchmark.

In constructing both measures, we apply standard data-cleaning procedures to mitigate the influence of reporting errors and extreme values. Specifically, observations with missing debt or interest expense are excluded, key balance sheet variables are winsorized at the 2.5th and 97.5th percentiles, and firm-year observations with implausibly high implied interest rates are removed. These filters ensure that the resulting cost-of-debt measures reflect economically meaningful borrowing costs.

C Theoretical Analysis: Central Banks and Collateral

We now extend the model to include the role of a central bank that intermediates liquidity provision through haircut-based collateralized lending. This extension preserves the core equilibrium logic of the model while enriching the structure of financial intermediation and allowing for the redistribution of default risk across institutions. It also lays the groundwork for policy design aimed at greening the financial system.

In the baseline model with $\theta = 1$, each firm k enters the period with a funding need equal to its liabilities $\mathcal{L}_k$. These liabilities must be financed up front at $t = 0$ in order to produce, and firms do not possess internal funds. In the benchmark case, banks directly supply this funding, bearing all the associated credit risk. In this extended setting, we introduce a central bank facility through which banks can refinance part or all of their lending using firm loans as collateral.

The key policy instrument of the central bank is the haircut parameter for sector k, denoted by $h_k \in [0, 1]$. The haircut determines the portion of a firm loan $\mathcal{L}_k$ that the central bank is willing to refinance. Green firms are totally re-financed, $h_g = 0$, while the haircut to brown firms h_b may be positive, and will be optimally chosen by the central bank to reduce systematic risk. A haircut of h implies that the central bank contributes $h \cdot \mathcal{L}_k$ toward the financing of firm k, while the remaining share, $(1-h) \cdot \mathcal{L}_k$, must be funded by the bank itself. Thus, the bank still performs the role of loan originator and underwriter, but the central bank becomes a co-financier.

At $t = 0$, each bank delivers the full loan amount $\mathcal{L}_k$ to the firm, as in the baseline model. However, it raises only part of this amount internally. The remainder is borrowed from the central bank against the collateral of the firm loan. This setup does not alter the real-side decisions of firms, they continue to operate under the same production technology and face the same shock vector. It does not affect the interest rate applied by banks, since their profit is

$$h_k \mathcal{I}_k \tag{77}$$

with $\mathcal{I}_k$ defined as in (17). Further, it does not affect the expectation of such profit, which remains zero. The set-up however affects the distribution of financial exposures across institutions.

At $t = 1$, firms realize their revenues $\mathcal{A}_k = p_k y_k^\eta$, which may fall short of liabilities due to productivity shocks. As in the baseline model, when firms cannot meet their obligations, lenders recover only a part of what they are owed.

The repayment flows to each intermediary are proportional to their share of the initial financing. Since the bank contributed a fraction $(1 - h_k)$ of the initial loan, it receives the same share of the recovery value. Similarly, the central bank receives the complementary fraction h_k. The bank's financial return from lending to firm k is thus given by:

$$h_k \mathcal{I}_k, \tag{78}$$

while the central bank's return is:

$$\mathcal{I}^{\mathrm{CB}} = \sum_{k \leq n_1} \mathcal{I}_k + \sum_{k > n_1} (1 - h_b) \mathcal{I}_k. \tag{79}$$

since the haircut is zero for green and h_b for brown sectors.

The total intermediation surplus from the loan therefore remains unchanged by the haircut. It merely reflects the overall gains or losses from financial intermediation and, for brown firms, is split between the central and the other banks.

In equilibrium, since expected profits from lending are zero for banks, they are zero also for the central bank:

$$\mathbb{E}[\mathcal{I}_k] = 0 \qquad \rightarrow \qquad \mathbb{E}[\mathcal{I}^{\mathrm{CB}}] = 0. \tag{80}$$

These conditions ensure that the full expected surplus from intermediation is passed through to households, in the form of higher consumption or lower interest rates, preserving the efficiency of the underlying allocation.

Importantly, the haircut does not distort firm-level incentives or resource allocation. Firms receive the full amount of their required funding $\mathcal{L}_k$ regardless of the value of h, and prices, wages, and production choices remain governed by the same equilibrium conditions. However, the haircut on brown firms plays a critical role in redistributing risk between intermediaries. A higher haircut shifts more risk to the central bank, reducing banks' exposure to firm default. Conversely, a lower haircut places more responsibility on banks to absorb losses, potentially leading to tighter credit conditions if banks become more risk-averse or capital-constrained.

Given its systemic role, the central bank may wish to set the haircut to manage its own risk exposure. One natural objective is to minimize the variance of its future cash flows, subject to maintaining efficiency and financial stability. The central bank's policy problem is thus:

$$\begin{gathered}\min_{h_b\in[0,1]} \mathrm{Var}\left(\mathcal{I}^{\mathrm{CB}}\right) \\ = \mathrm{Var}\left(\sum_{k\le n_1}\mathcal{I}_k\right) + \min_{h_b\in[0,1]}(1-h_b)^2\mathrm{Var}\left(\sum_{k>n_1}\mathcal{I}_k\right) + \\ +2(1-h_b)\mathrm{Covar}\left(\sum_{k\le n_1}\mathcal{I}_k, \sum_{k>n_1}\mathcal{I}_k\right)\end{gathered}$$

where the summation reflects aggregate exposure across the economy.

The necessary and sufficient condition for this minimization handles the following optimal value for the haircut

$$1-h_b = -\frac{\mathrm{Covar}\left(\sum_{k\le n_1}\mathcal{I}_k, \sum_{k>n_1}\mathcal{I}_k\right)}{\mathrm{Var}\left(\sum_{k>n_1}\mathcal{I}_k\right)} \tag{81}$$

where, since we are in the baseline model with $\theta = 1$, $\mathcal{I}_k = -\frac{\epsilon_k}{1+r_k} + \min(\tau_k, \epsilon_k)$.

Since the central bank profits inherit the volatility and correlation structure of the underlying shocks, including network amplification in τ, ϵ, this objective captures the central bank's desire to minimize macro-financial risk. This formulation also provides a natural foundation for green-oriented credit policy.