Financing Innovation: Leverage, Default, and Reallocation

Ufuk Akcigit Harun Alp Ege Y. Ercan*

March 26, 2026

Abstract

Growth-oriented firms borrow more, innovate more, and default more. We build a continuous-time Schumpeterian growth model in which heterogeneous firms differ in size, leverage, and innovation potential, and banks set firm-specific interest rates and debt limits. We calibrate the model to U.S. firm-level data from FR Y-14 and Compustat, matching joint distributions of leverage, spreads, default, and growth. In the calibrated economy, high-innovation firms hold over half of aggregate debt yet pay lower average spreads than low-innovation firms at comparable leverage, because their higher continuation values reduce default risk. Equalizing borrowing costs across firms raises creative destruction by 15 percent, indicating that the misallocation of credit across heterogeneous firms is quantitatively important for growth. We then analyze a temporary crisis that raises credit spreads. The ensuing bankruptcy wave reallocates product lines toward firms with stronger balance sheets and higher growth potential, generating a post-crisis overshooting of productivity growth. Faster bankruptcy resolution and reallocation toward high-innovation buyers amplify this effect.

JEL Codes: G3, O4

Keywords: Productivity, Crisis, Transition Dynamics, Debt, Financial Frictions, Default, Creative Destruction

*Author affiliations. Akcigit: University of Chicago, Department of Economics, NBER, & CEPR (uakcigit@uchicago.edu); Alp: Federal Reserve Board of Governors (harun.alp@frb.gov); Ercan: Stanford University, Graduate School of Business (egeercan@stanford.edu). The views in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or any other person associated with the Federal Reserve System.

1 Introduction

Firms' financial positions and capital structures play a central role in shaping investment, innovation, and survival. Differences in leverage and access to external finance influence which firms expand, which retrench, and which exit. When borrowing is distorted by financial frictions—stemming from default risk, limited pledgeability of assets, or regulatory and institutional constraints—the cost of funds can differ sharply across firms. These wedges have first-order implications for the allocation of resources across firms and for aggregate productivity. Growth-oriented, innovation-intensive firms may be constrained by weak balance sheets and high borrowing costs, while mature, low-growth incumbents with ample collateral can continue to roll over legacy debt. The resulting misallocation is present in normal times, but it becomes particularly salient during crises, when credit conditions tighten abruptly, interest-rate spreads jump, and bankruptcy and restructuring reallocate technologies and financial slack across firms. Whether this process primarily cleanses the economy by shifting resources toward high-innovation firms with sound balance sheets, or instead entrenches low-growth incumbents and suppresses creative destruction, is an open quantitative question with important implications for the depth of downturns and the strength of recoveries.

This paper develops a quantitative framework that embeds firm-level financial frictions into an endogenous growth model with creative destruction and rich firm heterogeneity. The framework is designed to study how borrowing costs, debt limits, and bankruptcy jointly shape productivity growth and selection in normal times, and how they govern the response of the economy to crises. We focus on two questions. First, how distortive are firm-specific borrowing costs and debt limits for innovation and creative destruction along the balanced growth path? Second, when a crisis tightens credit, does the ensuing wave of default and restructuring mainly destroy productive capacity or instead reallocate technologies and financial slack toward firms with stronger balance sheets and higher growth potential?

Our starting point is a Schumpeterian growth model in which incumbent firms produce differentiated intermediate goods and expand by innovating on random product lines, in the spirit of Klette and Kortum (2004) and Akcigit and Kerr (2018). Firms are heterogeneous along three dimensions that are central for both innovation and finance: their current scale, an innovativeness type capturing long-run growth potential, and their balance-sheet position, including accumulated debt and borrowing terms. High-type

firms are more efficient innovators and account for a disproportionate share of aggregate growth, while low-type firms are relatively stagnant. This persistent heterogeneity allows crises and policies to affect not only the average level of investment but also the composition of active firms and thus the path of future productivity.

Credit is provided by a representative bank that observes firms' states and offers short-term risky debt contracts with endogenous spreads and debt limits. For each borrower, the bank trades off higher interest margins against default risk and expected recovery values in bankruptcy. Debt is tax-advantaged but interacts with equity adjustment costs, so firms jointly choose payout policies, leverage, and innovation while managing the evolution of their balance sheets. Firms can exit voluntarily or default and enter a bankruptcy process in which their product lines are temporarily controlled by the bank and eventually sold at a discount to new owners. This environment generates firm-specific credit spreads, leverage distributions, default, and selection jointly with innovation and aggregate growth.

We discipline the model using detailed micro data on U.S. non-financial firms. We combine supervisory loan-level data, which provide information on bank–firm exposures, interest-rate spreads, and default, with Compustat data on public firms' payouts. We construct moments describing the joint distribution of leverage, spreads, growth, and exit across size and leverage groups, as well as entry rates, exit rates, and aggregate productivity growth, and choose parameters to match these moments. Particular emphasis is placed on reproducing the empirical patterns that growth and leverage are positively correlated in the cross-section, that interest-rate spreads increase in leverage and are higher for small firms, and that default is concentrated among highly levered firms.

The calibrated model displays three central features in steady state. First, high-innovation firms are substantially more levered and account for a large share of both growth and debt, yet they pay lower average spreads than low-innovation firms at comparable debt levels because their higher continuation values reduce default risk. Conditional on innovation type, however, leverage is growth-reducing in the model: more indebted firms invest less in innovation because high spreads and a tighter default region increase the shadow cost of internal funds. The positive correlation between leverage and growth observed in the data therefore reflects selection rather than a causal effect of debt on innovation. Second, firm-specific borrowing costs generate misallocation of innovation across firms. Equalizing spreads across all debtors raises creative destruction by about 15 percent, indicating that the cross-sectional dispersion in borrowing costs is quantitatively

important for aggregate growth. Third, bankruptcy and reallocation play a modest but non-negligible role in steady state. Endogenous defaults reallocate about two percent of GDP per unit time, and roughly 60 percent of these reallocations involve high-innovation firms, highlighting that financial distress affects precisely the firms that are central for long-run growth.

We then use the model to analyze a financial crisis that takes the form of a temporary increase in intermediation costs, which shows up as a spike in credit spreads, especially for firms close to their debt limits. The shock leads to an immediate fall in innovation and growth and to a wave of defaults. The key quantitative insight is that the same shock that destroys some productive capacity also generates a bankruptcy wave that reallocates assets and balance-sheet capacity toward firms with better financial health and higher growth potential. In our baseline calibration, this active reallocation is quantitatively important: it accounts for a large fraction of the medium-run rebound in productivity growth after the crisis. Absent this channel, the crisis has more persistent adverse effects, and the economy returns to its pre-crisis growth path more slowly and without overshooting in terms of productivity growth.

We show that the strength of this effect depends critically on institutional features of the bankruptcy process. Faster resolution of bankruptcy cases improves recoveries and reduces the time that technologies remain idle, allowing high-innovation buyers to start investing on acquired lines more quickly. Likewise, directing bankrupt assets toward buyers with higher innovation potential—such as entrants or incumbents with strong balance sheets—strengthens the recovery. In a counterfactual where bankrupt assets are always sold to entrants, who are more likely to be high-innovation firms and start with clean balance sheets, the crisis generates a substantially larger and more persistent improvement in post-crisis productivity.

Related Literature. Our paper is at the intersection of firm dynamics, financial frictions, Schumpeterian growth, and the macroeconomics of crises. On the firm dynamics and finance side, a large literature studies how borrowing constraints and heterogeneous access to finance affect investment, capital accumulation, and aggregate outcomes; see Cooley and Quadrini (2001), Jermann and Quadrini (2012), Miao (2005), and Li et al. (2016). Relative to this work, we embed rich, state-contingent credit terms and default into a model where firms endogenously choose innovation and where growth is driven by creative destruction. Our focus on the joint distribution of leverage, innovation potential, and

default risk connects this literature to empirical work on misallocation and dispersion in marginal products of capital.

Our growth block builds on Schumpeterian models with firm heterogeneity and quality ladder dynamics (Klette and Kortum, 2004; Akcigit and Kerr, 2018; Acemoglu et al., 2018). These papers emphasize the role of creative destruction in reallocating market shares across firms with different innovation intensities. We extend this framework by allowing firms to borrow and default, and by modeling bankruptcy explicitly as a mechanism for reallocating technologies and balance sheets. Papers such as Malamud and Zucchi (2019), Laeven et al. (2015), and Chatterjee and Eyigungor (2020) also study the interaction between finance and innovation, emphasizing liquidity constraints, credit screening, and low interest rate environments. The paper most closely related to ours in terms of modeling is Geelen et al. (2020), who compare different financial frictions in an endogenous growth setting. We differ by endogenizing firm-specific spreads and debt limits, modeling the bankruptcy process and post-bankruptcy reallocation explicitly, and analyzing transition dynamics during crises.

Our analysis speaks to the long-standing debate on whether recessions are cleansing or sullying. The classic view emphasizes that recessions may raise productivity by forcing inefficient units to exit, freeing resources for more productive firms (Saint-Paul, 1993; Aghion and Saint-Paul, 1998; Cooper and Haltiwanger, 1993; Caballero and Hammour, 1994). Subsequent work stresses that financial and labor market frictions can overturn this mechanism and lead crises to depress productivity through misallocation and stifled creative destruction (Barlevy, 2002; Aghion et al., 2010; Caballero and Hammour, 2005; Cerra and Saxena, 2008; Foster et al., 2016). We contribute by showing, in a model calibrated to micro data on leverage and spreads, how bankruptcy-driven reallocation in a crisis can be sharply pro-growth if it directs assets toward firms with superior innovation capacity and healthier balance sheets.

Finally, our work relates to macro papers that analyze recessions and credit shocks in models with firm heterogeneity and financial frictions, such as Gopinath et al. (2017) and Osotimehin and Pappada (2017), and to studies that revisit the long-term growth consequences of stabilization policies in turbulent times, as in Blackburn (1999). We differ from these contributions by embedding endogenous innovation and debtor-specific credit terms into a Schumpeterian framework and by decomposing post-crisis recoveries into the contribution of continuing incumbents and that of reallocated product lines emerging from bankruptcy. In addition, by explicitly modeling recoveries and the speed

of resolution, our framework connects macroeconomic dynamics to institutional features of bankruptcy law and distressed asset markets.

The remainder of the paper is organized as follows. Section 2 presents the model. Section 3 describes the data and calibration strategy and discusses identification. Section 4 analyzes the steady-state equilibrium and the role of financial frictions for innovation, selection, and misallocation. Section 5 studies the transition dynamics of a financial crisis and the role of bankruptcy and reallocation. Section 6 concludes.

2 Model

We consider a firm-based model of endogenous technical change in the spirit of Klette and Kortum (2004) and Akcigit and Kerr (2018), augmented with three ingredients: (i) firm-level borrowing with endogenous state-contingent spreads and debt limits, (ii) a bankruptcy process that reallocates technologies across firms, and (iii) persistent heterogeneity in innovation productivity that interacts with balance-sheet position. We now describe the environment and characterize the stationary balanced growth equilibrium.

2.1 Household and Final Good Technology

We consider a continuous time economy, where a representative household has CRRA preferences over consumption:

$$U_0 = \int_0^\infty e^{-\rho t} \frac{C_t^{1-\vartheta} - 1}{1-\vartheta} dt,$$

The household maximizes its utility subject to the flow budget constraint

$$\dot{D}(t) + \dot{A}(t) + C(t) \leq r_f(t)D(t) + r_f(t)A(t) + W(t)L^S,$$

where $A(t) = \int_{i\in\mathcal{N}(t)} V(t;i)di$ is the aggregate equity stake in firms, $\mathcal{N}(t)$ is the set of active firms, $V(t;i)$ is the ith firm's value, $r_f(t)$ is the risk-free interest rate, $D(t) = \int_{i\in\mathcal{N}(t)} D(t;i)\, di$ is the aggregate debt level, $D(t;i)$ is the ith firm's debt, L^S is the household's measure of fixed labor, $W(t)$ is the wage rate, and the usual no-Ponzi conditions $\int_0^\infty e^{-r_f(t)t}A(t)dt \geq 0$ and $\int_0^\infty e^{-r_f(t)t}D(t)dt \geq 0$. Household's maximization implies the standard Euler equa-

tion, determining the risk-free interest rate, r_f:

$$\frac{\dot{C}}{C} = \frac{r_f - \rho}{\vartheta}. \tag{1}$$

The final good, Y, is produced by aggregating intermediate good varieties, y_i, according to the following technology:

$$\ln Y = \int_{i \in \Phi} \ln y_i di,$$

where Φ is the set of active intermediate good product lines, the measure of which is normalized to one. The final good has two uses: (1) it is consumed by the household and (2) expended by intermediate good producers to invest in productivity enhancement.

The final good sector's demand schedule for each intermediate good variety is

$$\arg \max_{\{y_i\}_{i \in \Phi}} P e^{\left(\int_{i \in \Phi} \ln y_i di\right)} - \int_{i \in \Phi} p_i y_i \, di,$$

where P is the price of the final good and p_i is the price of variety i. The above optimization results in unit-elastic demand for each variety:

$$y_i = \frac{Y}{p_i},$$

where we normalize the price of the final good to one.

2.2 Intermediate Good Production

Intermediate good i is produced by the monopolist who has the most advanced technology in that product line. A single monopolist can own multiple product lines and can produce multiple intermediate goods simultaneously. Each intermediate good is produced with a linear technology based on the intermediate good producer's *variety-specific* productivity, $q_i > 0$, and labor, $l_i \geq 0$,

$$y_i = q_i l_i.$$

For each variety, firms compete à la Bertrand to undercut their competitors' prices. By limit pricing, only the most productive producer gets to produce the variety. The most

productive firm is $\lambda > 1$ times as productive as its competitors, where λ is the improvement step size. Variety i's equilibrium price is

$$p_i = \lambda \frac{W}{q_i}, \tag{2}$$

where W is the wage rate. The equilibrium production is

$$y_i = \frac{q_i Y}{\lambda W}, \tag{3}$$

and the intermediate good producer's profit from producing variety i is

$$\pi = \frac{q_i Y}{\lambda W} \lambda \frac{W}{q_i} - \frac{q_i Y}{\lambda W} \frac{1}{q_i} W = \frac{\lambda - 1}{\lambda} Y.$$

Lastly, note that given a fixed labor supply, L^S, the labor market clears with the wage rate

$$W = \frac{||\Phi||}{L^S \lambda} Y. \tag{4}$$

2.3 Firm Heterogeneity and Dynamics

Incumbent intermediate good producers are heterogeneous in terms of the number of intermediate goods they produce ($n \in \mathbb{N}$), debt stock ($D \in \mathbb{R}$), and innovativeness type ($k \in \{H, L\}$). A firm's type represents its innovative productivity, and high-type firms ($k = H$) are more productive at innovating than low-type firms ($k = L$). The low-type state is absorbing, and high-types exogenously transition into it with some arrival rate, $\phi_{H \to L} > 0$.

Productivity enhancement is modeled as an improvement over a product line's pre-existing productivity by the improvement step size $\lambda > 1$. Innovation is undirected meaning that the firm does not target specific intermediate varieties but instead achieves an improvement over a random variety each time.

When a firm successfully improves on the technology of a variety, it captures that product line by undercutting the prices of the previous producer. This process can be described as "creative destruction" ($\tau \geq 0$) as the turnover of producers is the driving force of endogenous growth.

An incumbent firm achieves a successful improvement at the Poisson rate, $X \geq 0$, which is a function of type-specific innovation productivity $\theta_k > 0$, size n, and investment $c \geq 0$,

$$X = \left(\theta_k \frac{c}{Y}\right)^{\frac{1}{\eta}} n^{\sigma},$$

where $\eta > 1$ is the inverse innovation elasticity, $k \in \{H, L\}$, and investments are made by expending final goods. High-type firms face an innovation productivity scale $\theta_H > 0$ while low types face the scale $\theta_L > 0$, where $\theta_H > \theta_L$. Rewriting the productivity flow expression, we get that the firm's investment cost function to achieve an productivity enhancement *intensity* of $x \geq 0$ (where $x = \frac{X}{n}$) is

$$c(x) = \frac{1}{\theta_k} x^{\eta} n^{\eta(1-\sigma)} Y.$$

In addition to expending funds to improve their productivity, incumbents also choose payouts. The present value of all future payouts determines the firm's value to its shareholder, the representative household which holds an aggregate basket of all firms' shares. Firms are subject to equity rigidities, modeled as quadratic equity adjustment costs similar to Jermann and Quadrini (2012). This cost structure creates a wedge between the value of the payout to the investor and the cost to the firm.

Let $p \in \mathbb{R}$ represent a firm's payout intensity (i.e., total payout over firm size). To issue pnY final goods in payouts, the firm has to expend

$$\varphi(p) = pnY + n\kappa p^2 Y,$$

where $\kappa > 0$ is the penalty scale. Note that the firm can choose a negative payout (i.e., an equity injection). However, the firm will have to tolerate potentially significant costs to raise the funds through negative payouts. We do not treat these costs as resource drains on the economy, and we return the expended equity adjustments costs to the household as a lump-sum payment.

Lastly, firms face a size invariant fixed cost, $\chi Y \geq 0$, which we do not treat as a resource drain on the economy. Then, the firm's cash flows from production, fixed cost,

productivity investments, and payout is

$$\pi n - \chi Y - c(x) - \varphi(p).$$

2.4 Firms and Debt

The firm has access to a defaultable short-term debt instrument, structured by the banking sector around each firm's heterogeneity and demand for debt. Namely, the firm is charged a firm-specific endogenous interest rate, $r(D, n, k) > 0$, and faces a firm-specific endogenous debt limit, $\bar{D}(n, k) \geq 0$, above which the firm cannot receive a loan.

The firm's debt evolves according to the following law of motion:

$$\dot{D}(x, p; D, n, k) = -\underbrace{[\pi n - \chi Y - c(x) - r(D, n, k)D]}_{\text{Before-Tax Net Income}}(1 - \tau_{tax}) + \varphi(p). \tag{5}$$

where $r(D, n, k) = r_f + 1\{D > 0\}\delta(D, n, k)$ and $\delta(D, n, k) \geq 0$ is the interest rate spread coming from the banking problem. Note that corporate tax confers tax shield status on interest expenses. By choosing its payout and investment, the firm also chooses its debt drift above. This drift is unconstrained when the firm is not at its debt limit.

At any moment, the firm can exercise its outside option by exiting the economy. Hence, the firm's value cannot drop below 0:

$$V(D, n, k) = \max\{0, V_{\text{cont.}}(D, n, k)\}, \tag{6}$$

where $V_{\text{cont.}}(D, n, k) \in \mathbb{R}$ is the continuation value. When the firm is encumbered with debt, it can exit through a bankruptcy. By limited liability, seeking bankruptcy protection allows the firm to avoid paying its debt obligations. But the bankruptcy process is still costly to the bank (e.g., retaining lawyers and paying court fees) meaning that the firm may choose to apply for bankruptcy even though it has positive equity value. Given firm characteristics (D, n, k), solving the exit decision, $B(D, n, k) \in \{0, 1\}$, boils down to solving for debt levels beyond which the firm's continuation value falls below 0.

Four types of shocks can lead to bankruptcy. First, the firm that experiences a death

shock (which occurs at the arrival rate $\varepsilon \geq 0$) is forced to shut down. Second, the firm that is solvent as a high-type can choose to default when it experiences a type transition shock (which occurs at the arrival rate $\phi_{H \to L} > 0$) that morphs the firm into a low-type firm. That is to say, there can exist firm characteristics (D, n, H) such that $B(D, n, H) = 0$ though $B(D, n, L) = 1$. Third, the firm exits when it is completely creatively destructed from the market (i.e., loses all of its product lines, where each product line is lost at the arrival rate $\tau \geq 0$). Lastly, the firm that loses a product line can declare bankruptcy despite operating other lines if it prefers defaulting to trying to repay its obligations. Specifically, there can exist firm characteristics (D, n, k) such that $B(D, n, k) = 0$ while $B(D, n-1, k) = 1$.

2.5 Banking

The banking sector is comprised of a representative bank that borrows funds from the household and extends loans to individual debtor firms. In this process, the bank observes debtor firms' heterogeneities and offers them a loan menu comprised of firm-specific interest rates, $r(D, n, k) \geq 0$ and debt limits, $\bar{D}(n, k) \geq 0$.[1]

In the credit market, the bank's primary instrument is setting interest rate spreads (i.e., the difference between the risky borrowing rate, $r(D, n, k)$, and the risk-free rate, r_f). The bank offers its loans at a premium to compensate itself for bankruptcy risks and miscellaneous banking costs.

Note that, at certain debt levels, firms lack any and all commitment to repay their loans. For example, if a very small firm receives a very large loan, the firm could optimally choose to immediately default on its obligations. Therefore, the bank would lose funds by extending a loan amount that large to a firm that small at any interest rate. In such circumstances, the bank utilizes debt limits to deny loans. Hence, in the economy, interest rate spreads and borrowing constraints are endogenously instituted by the bank.

Consider the case of a k-type firm with size n applying for loan amount $D > 0$. If the firm does not suffer from a commitment problem, the bank extends the loan at a firm-specific interest rate spread, $\delta(D, n, k) \geq 0$.

[1] The underlying modeling choice is that the bank sets interest rates for each debt level. At debt levels for which the firm cannot credibly repay its obligations, the bank charges infinite interest rate, thus avoiding extending the loan. That is, the bank effectively chooses debt limits above which it is unwilling to extend credit. This treatment is similar to Bornstein (2020).

The bank funds this loan by borrowing from the household at the risk-free rate, r_f. The bank is able to borrow at the risk-free rate because it extends loans to a continuum of firms, perfectly diversifying its credit risks. Hence, the bank's total gains and losses are known with certainty, meaning that it can structure its loans such that it always honors its obligations to the household.

The debtor firm defaults with some flow rate. When a bankruptcy occurs, the firm has some size-dependent liquidation value, potentially allowing the bank to recover some funds. Let $E[\Phi_D(D, n, k)] \geq 0$ represent the expected ratio of unrecovered funds over the loan's principal. Capturing other sources of banking costs, some amount, ν_0, is expended by the bank for every unit of loan it extends. Then, when granting a loan, D, to the the firm characterized by (n, k), the bank's profit function is

$$\pi_B(D, n, k) = \underbrace{D(1 + r_f + \delta(D, n, k))}_{\text{Debtor Firm's Obligations}} - \underbrace{D\mathbb{E}[\Phi_D(D, n, k)]}_{\substack{\text{Expected Losses} \\ \text{from Bankruptcies}}} - \underbrace{D(1 + r_f)}_{\substack{\text{Bank's} \\ \text{Obligations}}} - \underbrace{D\nu_0}_{\substack{\text{Banking} \\ \text{Costs}}} . \tag{7}$$

The bank is risk-neutral and prices its loans competitively. Hence, the equilibrium spread is determined as the rate that sets the bank's profits to zero:

$$\delta^*(D, n, k) = \underbrace{\nu_0}_{\substack{\text{Banking} \\ \text{Costs}}} + \underbrace{\mathbb{E}\left[\Phi_D(D, n, k)\right]}_{\text{Expected Loss Ratio}}, \tag{8}$$

Observe that the equilibrium credit spread function, when charged over a continuum of firms, ensures that the bank's income from firms less its other costs exactly equals its obligations to the household.

In order to deny loans that would lead to the firm's defaulting with certainty, the bank sets debt limits as follows:

$$\bar{D}(n, k) = \max\{D : V_{Firm}(D, n, k) \geq o_{close} n\} \tag{9}$$

o_{close} is the cost of closing down the firm, triggering the firm's liquidation before it hits non-positive continuation value.

Importantly, default rates and recoveries are not functions of the firm's *current* decisions. Instead, they are a function of the culmination of *past* actions and shocks (i.e, the

firm's current state). Therefore, the interest rate schedule is characterized completely by setting interest rates for possible firm heterogeneities (D, n, k) without needing to price the firm's current choices of x and p.

2.6 Bankruptcy

When a firm defaults, it enters the bankruptcy process. During this process, income generated by the bankrupt firm's product lines are paid to the bank. At poisson arrival $\psi > 0$, those product lines are sold to a buyer and the old firm ceases. This delay captures the court process and search frictions present in the real economy. We model bankruptcy as a stylized fusion of Chapter 11 (i.e., reorganization) and Chapter 7 (i.e., liquidation) bankruptcies. Sales occur as a result of a negotiation between the buyer and seller (i.e., the bank). The value of the resold firm depends on the type of the buyer who is a high type with probability $\alpha \in [0, 1]$. Both agents know the buyer's type and negotiate accordingly. The sale is facilitated by a new loan given to the buyer to purchase the firm.

In the negotiation, the bank's threat point is zero and its leverage is $(1 - \chi) \in [0, 1]$. The buyer's threat point is the debt limit of a firm of the given size and an owner with the given type. Intuitively, the debt limit marks the debt level at which the firm has zero equity value. The buyer would be willing to finance the purchase (weakly) up to this point because she would gain non-negative equity value without any costs to herself. By Nash Bargaining, the equilibrium sale price of a firm with size $n > 0$ and buyer with type k is

$$\bar{D}(n, k)\mu$$

The bank does not know the future buyer's type when it extends credit to the debtor firm. Let $m_{B,H} \equiv \alpha$ denote the probability that the buyer is a high type. Bank's expected recoveries if a firm with size $n > 0$ is sold is

$$Q(n) = (1 - \xi) \left[m_{B,H} \bar{D}(n, k = H) + (1 - m_{B,H}) \bar{D}(n, k = L) \right].$$

Moreover, the bank does not know how quickly it will be able to resell the firm, which continues to be subjected to creative destruction during bankruptcy. Calculating firm's

value to the bank involves solving the following value function:

$$r_f R_L(n) = \max\Big\{0, \pi n - \chi + \tau n \left[R_L(n-1) - R_L(n)\right] + \psi \left[Q(n) - R_L(n)\right]\Big\}. \tag{10}$$

Importantly, this market structure ensures that endogenous defaults do not cause a reduction in contemporaneous output. Alternative to the expected recovery in line (10), the bank can claim its outside option $\mu_0 Y \geq 0$, which captures recovery value unaccounted in our theoretical model.

The bank prices its loan conscious of the different liquidation values under different bankruptcy scenarios. Let $1_\phi = 1_{k=H} B(D, n, L)$ be the indicator variable representing whether firm characterized by (D, n, k) would declare bankruptcy when transitioned to a low-type. Let $1_\tau = B(D, n-1, k)$ be the indicator variable representing whether the firm would default if it shrinks by a size. Then, the bank faces an expected loss ratio equaling

$$\mathbb{E}[\Phi_D(D,n,k)] = \underbrace{(\varepsilon + \phi_{H\to L} 1_\phi)}_{\substack{\text{Flow Rate}\\ \text{of Size } n\\ \text{Default}}} \underbrace{\left[\frac{D - \max\{\mu_0, R_L(n)\}}{D}\right]^+}_{\substack{\text{Ratio of}\\ \text{Unrecovered}\\ \text{Debt at Size } n}} + \underbrace{\tau n 1_\tau}_{\substack{\text{Flow Rate}\\ \text{of Size } n-1\\ \text{Default}}} \underbrace{\left[\frac{D - \max\{\mu_0, R_L(n-1)\}}{D}\right]^+}_{\substack{\text{Ratio of}\\ \text{Unrecovered}\\ \text{Debt at Size } n-1}},$$

where $[]^+$ stands for the positive part of the expression in the brackets. That is, the bank does not recover more funds in the bankruptcy than it lent

2.7 Entry and Exit

As we focus on incumbent financial frictions, we model the entrant's problem more parsimoniously. Potential entrants are homogeneous until entry and have a mass of one. They invest towards their productivity to capture a product line and enter the market, at which time they enter without debt/savings (i.e., $D = 0$), hold one product line, and draw a growth type. Each entrant has probability $\alpha \in [0, 1]$ of becoming a high type and probability $1 - \alpha$ of becoming a low type.

The entrant's innovation dynamics is similar to the incumbent's. The following is the

entrant innovation cost function:

$$c_E(z) = \frac{1}{\theta_E} z^{\eta} Y,$$

where z is the arrival rate of successful innovations and $\theta_E > 0$ is the innovation productivity scale.

By choosing only their innovation intensity as follows, entrants maximize their expected value:

$$z\left[\alpha V(0,1,H) + (1-\alpha)V(0,1,L)\right] - c_E(z). \tag{11}$$

Potential entrant innovation presents the last source of incumbent exit in our model.

2.8 Firm Value

Recursively, firm value is

$$\begin{aligned} & r_f V(D,n,k) - \partial_t[V(D,n,k)] \\ & = \max\Bigg\{0,\ \max_{x,p}\Bigg\{pnY + \frac{\partial V}{\partial D}(D,n,k)\dot{D}(x,p;D,n,k) - \varepsilon V(D,n,k) \\ & \qquad + xn\left[V(D,n+1,k) - V(D,n,k)\right] + \tau n\left[V(D,n-1,k) - V(D,n,k)\right] \\ & \qquad + 1_{k=H}\phi_{H\to L}\left[V(D,n,L) - V(D,n,H)\right]\Bigg\}\Bigg\} \end{aligned} \tag{12}$$

subject to $D \leq \bar{D}(n,k)$ and $\dot{D}$ as in (5).

The outer maximization is the firm's exit decision. The inner maximization is the firm's investment and payout decisions if it continues to operate. The first term accounts for the contribution of payouts to firm value. The second term accounts for how firm value changes as a result the change in debt. The third term accounts for the exogenous death shock's effect on firm value. The fourth term accounts for the value of innovation, and the fifth term captures the value lost due to creative destruction. The sixth line accounts for value destruction when high types transition into low types.

2.9 Aggregate Growth and Welfare

Proposition 1. *The growth rate of the economy is equal to*

$$g = \ln(\lambda)\tau \tag{13}$$

Definition A stationary equilibrium of this economy is a tuple

$$\{y_i, p_i, W, V, B, x, p, z, \delta, \bar{D}, R_L, g, r_f\}$$

such that y_i and p_i maximize the ith product line's profits as in (3) and (2), W clears the labor market as in (4), V is the incumbent firm value function as in (12) subject to the debt limit $\bar{D}(n,k)$ and the debt drift as in (5), $B(D,n,k)$ underpinning voluntary market exits in accordance with (6), $x(D,n,k)$ and $p(D,n,k)$ optimizing incumbent firm value $V(D,n,k)$, z maximizing entrant value as in (11), $\delta(D,n,k)$ maximizing the bank's expected profits in (7) as in (8), $\bar{D}$ ensuring that a loan is extended only when $\delta(D,n,k)$ has a finite solution as in (9), R_L captures the collateral value as in (10), g is given by (13), and r_f is determined in accordance with the Euler equation in (1).

3 Data and Calibration

We now describe the data used to discipline the model and our calibration strategy. The goal is to reproduce key cross-sectional patterns in leverage, credit spreads, growth, and default, as well as aggregate growth, while preserving a transparent mapping from parameters to moments.

3.1 Data Sources and Variable Construction

All moments except one are constructed from the supervisory FR Y-14 data set, spanning 2013–2019. FR Y-14 collects confidential loan-level data from large bank holding companies subject to Federal Reserve stress testing. The corporate credit module reports, at annual frequency, firm-level information on outstanding balances, interest rates, contractual spreads, internal risk ratings, and realized default events. We use the panel of

non-financial firms with at least one term loan or revolving credit facility. Because these data originate from regulatory reports rather than voluntary filings, coverage is broad and includes many privately held firms that are not observed in standard commercial databases.

We complement FR Y-14 with Compustat data on public firms to measure payout behavior. FR Y-14 contains detailed information on debt but not on equity repurchases or dividends. Because payout behavior is particularly relevant for large firms, we restrict the Compustat sample to the top decile of the firm-size distribution, as measured by sales, and compute payout moments in that subsample. To minimize discrepancies in definitions across data sets, we harmonize accounting variables such as sales and debt wherever possible.

We construct several key variables that are used both in calibration and in the descriptive analysis. Firm size is measured by total sales. We define leverage as the ratio of interest-bearing debt to sales. For firms with multiple loans in FR Y-14, we aggregate loan-level information to the firm-year level using outstanding balances as weights. The interest-rate spread is defined as the contractual interest rate minus the risk-free rate appropriate to the maturity of the loan. We focus on spreads over short-term benchmarks and average across loans for each firm-year observation. Default is a firm-year indicator equal to one if the firm experiences a default event on any of its loans during the year. To construct growth, we use the year-on-year log change in sales, trimming the top and bottom one percent of the distribution to reduce the influence of outliers.

We drop financial firms (NAICS 52) and utilities (NAICS 22), as their capital structure decisions are subject to different regulatory constraints. We also exclude firms with missing sales or leverage. Our final sample contains more than 600,000 firm-year observations and about 190,000 distinct firms. These firms account for roughly 80 percent of U.S. aggregate output, whereas public firms in Compustat account for about 40 percent. The overlap between the two sources is concentrated in the upper part of the firm-size distribution, which motivates our use of Compustat only for payout moments.

Measurement issues are important for mapping the model to the data. Our notion of leverage corresponds to debt relative to total revenues, whereas some of the empirical corporate finance literature uses book or market leverage ratios based on assets or equity. We adopt the debt-to-sales measure because it aligns more closely with the model, in which debt and profits both scale with the number of product lines, and because it is

consistently available in FR Y-14. We also interpret the model's default events as corresponding to the regulatory notion of default used in FR Y-14, which includes both missed payments and distressed restructurings.

3.2 Stylized Facts

The calibration targets are chosen to capture a set of stylized facts about firm finance and growth that motivate our modeling choices.

First, leverage and interest-rate spreads are jointly heterogeneous across firms. In FR Y-14, larger firms are more likely to have bank debt, but among indebted firms, smaller firms are more levered on average. Conditional on firm size, higher leverage is associated with higher spreads, and spreads are systematically higher for small than for large firms. This pattern is consistent with an endogenous default risk channel operating through both size and debt, and it motivates our emphasis on firm-specific debt limits and state-contingent spreads.

Second, growth and leverage are positively correlated in the cross-section. Firms in the upper half of the leverage distribution grow faster on average than firms in the lower half, both among small firms and among large firms. At the same time, high-leverage firms default more often. The joint distribution of leverage, growth, and default suggests that high-growth firms endogenously select into higher leverage, but also that high leverage increases the risk of distress when conditions deteriorate. This combination of selection and fragility is central for our quantitative analysis of crises.

Third, spreads are more volatile for highly levered firms and respond more strongly to macro conditions. Comparing years with elevated aggregate stress to normal times, spreads on loans to high-leverage firms rise more than spreads on loans to low-leverage firms. This non-linearity motivates our focus on an intermediation cost shock that disproportionally tightens borrowing conditions for firms near their debt limits.

Fourth, payouts are pervasive among large public firms. In the Compustat sample, about 84 percent of large firms pay out in a given year, and the mean payout conditional on paying is about 6 percent of sales. At the same time, payout and investment decisions are not frictionless: equity issues are rare and often concentrated in a subset of young firms, while seasoned firms adjust payouts gradually over time. These patterns motivate

the inclusion of equity adjustment costs and inform the calibration of the corresponding parameters.

Finally, entry and exit are substantial. Entry rates are on the order of two percent per year, and exit rates are non-negligible even among larger firms. Although not all exit events correspond to bankruptcy, the persistence of entry and exit suggests that creative destruction and selection are quantitatively important. In the model, we will discipline both the rate at which new firms enter and the exit behavior of incumbents by moments from these distributions.

3.3 Parameterization and Moments

We distinguish between externally and internally calibrated parameters. The externally calibrated parameters are

$$\{\rho,\ \vartheta,\ \eta,\ \tau_{tax},\ \psi\},$$

and the internally calibrated parameters are

$$\{\kappa,\ \nu_0,\ \xi,\ \mu_0,\ \chi,\ o_{close},\ \theta_H,\ \theta_L,\ \theta_E,\ \lambda,\ \alpha,\ \phi_{H\to L}\}.$$

The external parameters are set to standard values: $\rho = 0.02$, $\vartheta = 2$, and $\eta = 2$, following Acemoglu et al. (2018); the corporate tax rate is $\tau_{tax} = 0.25$, in between the calibrations of Geelen et al. (2020) and Jermann and Quadrini (2012); and the default resolution rate $\psi = 0.46$ is chosen to match the average bankruptcy duration reported by Bris et al. (2006).

The internal parameters are identified via simulated method of moments. We choose them to minimize a distance between model moments model(i) and data moments data(i),

$$\sum_{i=1}^{16} \frac{w(i)|\text{model}(i) - \text{data}(i)|}{\frac{1}{2}|\text{model}(i)| + \frac{1}{2}|\text{data}(i)|}, \tag{14}$$

where $w(i)$ is a weight. The sixteen moments include interest-rate spreads and leverage by size and leverage groups, the distribution of payouts, growth rates by size and leverage, entry and exit rates, and aggregate growth. The weighting scheme penalizes proportional deviations, so that moments that are close to zero receive lower weight.

We explore the parameter space in three stages. First, we generate candidate param-

eter vectors via Sobol sequences and solve the model for each. This step provides a coarse mapping from parameters to moments. Second, we apply simulated annealing to the most promising candidates, allowing for global exploration of the objective function. Third, we perform local optimization starting from the best annealed parameter sets. For each parameter vector, we solve the equilibrium by finite differences in the spirit of Achdou et al. (2017), with modifications to handle endogenous debt limits and spreads. The computational strategy is summarized in Appendix B.

In addition to the targeted moments, we use several non-targeted moments as informal checks on the model. For example, we compare the implied distribution of interest-coverage ratios, the correlation between spreads and ex post default, and the age distribution of borrowers to their empirical counterparts. The model reproduces the broad shape of these distributions, although it does not target them directly. These additional comparisons give us confidence that the parameterization is not over-fitting a narrow set of statistics.

3.4 Identification

Several groups of parameters are tightly linked to specific sets of moments.

On the financial side, the equity adjustment cost parameter κ governs how aggressively firms adjust payouts in response to shocks. We identify κ by matching the fraction of firms with positive payouts and the mean payout conditional on paying out, using Compustat for the top decile of firms by size. Intuitively, a higher κ makes equity issuance and payout adjustments more costly, reducing the fraction of firms that pay out and lowering the conditional payout rate. Fixed operating costs χ reduce effective cash flow, especially for small firms, and are crucial for their debt capacity. We therefore use leverage among small firms in FR Y-14 to discipline χ. Increasing χ reduces the debt capacity of small firms relative to large firms, steepening the size-leverage gradient.

The negotiation parameter ξ and the closure threshold o_{close} determine the extent to which banks recover in bankruptcy and thus shape debt limits, particularly for large firms. A higher ξ raises recoveries, allowing banks to extend more credit to large firms without increasing expected losses. We identify these parameters using leverage of large, highly levered firms. The fixed recovery value μ_0 affects the recovery of small bankrupt firms and influences spreads and leverage at the bottom of the size distribution. If μ_0 is

too low, small firms face very tight debt limits and cannot achieve the leverage observed in the data. The banking cost ν_0 is pinned down mainly by average spreads among large, relatively safe firms, where default risk plays a smaller role. Raising ν_0 shifts spreads upward for all firms, while leaving the cross-sectional dispersion driven by $\mathbb{E}[\Phi_D]$ intact.

On the innovation side, the high- and low-type innovation productivities θ_H and θ_L, together with the share of high-type entrants α and the transition rate $\phi_{H\to L}$, determine life-cycle growth rates across size and leverage groups and the composition of types along the size distribution. If θ_H is too high relative to θ_L, the model produces excessively fast growth for large firms and too much concentration of innovation among high types, making it difficult to match growth rates in high-leverage groups without overshooting aggregate growth. Lowering α reduces the share of high types among small firms and shifts innovation toward incumbents, affecting entry and exit rates. The innovation productivity of entrants θ_E shapes the entry margin and the share of small firms; we identify it from the entry rate and the exit rates of small incumbents.

Finally, the exogenous death rate ε affects exit almost uniformly across the size distribution. If ε is too low, exit is dominated by endogenous bankruptcy and creative destruction, and the model struggles to match exit rates among large firms that rarely default. We match exit rates of large incumbents, jointly with the exit of small firms, to discipline ε. The step size λ pins down aggregate growth given the rate of creative destruction and thus is calibrated to match aggregate productivity growth. A larger λ amplifies the contribution of each reallocation event to growth, allowing a given rate of creative destruction to sustain higher g.

Table 1 summarizes the externally calibrated parameters.

Table 1. *External Calibration*

#	Parameter	Description	Value
1	ρ	Discount rate	0.02
2	ϑ	Inverse intertemporal elasticity of substitution	2.00
3	η	Inverse innovation elasticity	2.00
4	τ_{tax}	Corporate tax rate	0.25
5	ψ	Bankruptcy resolution rate	0.46

4 Steady-State Results

4.1 Calibrated Parameters

Table 2 reports the internally calibrated parameters. It is useful to summarize their economic implications.

Table 2. *Internal Calibration*

#	Parameter	Description	Value
1	κ	Payout adjustment cost scale	2.124
2	χ	Fixed operating cost (in units of Y)	0.225
3	μ_0	Fixed recovery value (in units of Y)	0.615
4	ξ	Bank's bargaining share	0.421
5	o_{close}	Closure threshold per product line	0.463
6	ν_0	Banking cost per unit of debt	0.013
7	θ_H	High-type innovation productivity	0.038
8	θ_L	Low-type innovation productivity	0.012
9	θ_E	Entrant innovation productivity	0.040
10	λ	Step size of innovations	1.400
11	α	Share of high types at entry	0.721
12	$\phi_{H\to L}$	Transition rate from high to low type	0.049

The payout adjustment cost κ implies that, if a firm were to pay out its entire profit flow in one instant, roughly one-third of that outlay would be absorbed by adjustment costs. This friction rationalizes the smoothness of payouts and makes equity financing particularly costly for small and financially constrained firms. The fixed operating cost χ is about 110 percent of the profit from one product line, rendering very small firms more fragile and giving rise to an economically meaningful extensive margin of exit. In the calibrated distribution, a non-trivial fraction of the smallest firms operate close to their breakeven condition.

On the financial side, the bargaining parameter ξ implies that, on average, assets sold out of bankruptcy trade at a discount of around 25 percent relative to their implied debt capacity. The fixed recovery value μ_0 corresponds to roughly 2.7 times the annual revenue of a firm with one product line, providing a floor on recoveries and allowing some very small bankrupt firms to generate substantial recoveries despite low resale values. The banking cost ν_0 accounts for about half of the average spread paid by small, highly levered firms and a smaller fraction for large firms. Together, these parameters produce

an equilibrium in which spreads reflect both expected losses and intermediation costs, as in equation (8).

On the innovation side, the calibrated innovation productivities imply that high-type firms are more than three times as innovative as low types. Entrant innovation productivity is comparable to that of high types, reflecting aggressive innovation at entry that is necessary to match empirically observed entry and exit rates. The step size $\lambda = 1.40$ implies that the profit rate from each product line is about 22 percent of sales. The share of high types at entry and the transition rate $\phi_{H \to L}$ imply that roughly three-quarters of entering firms are high type, but an individual high-type firm has an expected high-type spell of about 3.7 years. This structure generates a non-degenerate mixture of types in the steady-state cross-section and allows the model to reproduce observed dispersion in growth rates.

4.2 Fit

Table 3 compares model and data moments. The model tracks the targeted moments closely and reproduces the qualitative patterns discussed in the data section.

The model reproduces the empirical pattern that large firms are more likely to have debt, but among indebted firms, small firms are more levered on average. Spreads across size and leverage groups are matched well for all cells except large, high-leverage firms, where the model overpredicts the spread (3.0 percent versus 1.5 percent in the data) and underpredicts leverage (0.49 versus 0.65). These two misses are related: the model generates steeper spread schedules for large firms than the data suggest, which in turn discourages those firms from levering up as aggressively. One possible explanation is that large, highly levered firms in the data enjoy relationship-lending benefits or reputational advantages that reduce their borrowing costs relative to the model's competitive pricing assumption. The growth moments show that the model captures the higher growth of highly levered firms within both size bins and the slower growth of large, low-leverage firms, while matching aggregate growth almost exactly. Non-targeted moments such as the distribution of default rates across leverage bins and the conditional correlation of spreads and default are also reasonably close to the data, suggesting that the model is not overfitting a narrow set of statistics.

Table 3. *Targeted Moments: Model vs Data*

#	Moment	Data	Model
1	Fraction of positive payouts (top size decile)	0.838	0.839
2	Mean payout conditional on positive	0.060	0.051
3	Interest spread (small size, low leverage)	0.014	0.013
4	Interest spread (small size, high leverage)	0.018	0.018
5	Interest spread (large size, low leverage)	0.013	0.013
6	Interest spread (large size, high leverage)	0.015	0.030
7	Entry rate	0.018	0.017
8	Leverage (small size, low leverage)	0.062	0.062
9	Leverage (small size, high leverage)	0.616	0.603
10	Leverage (large size, low leverage)	0.070	0.091
11	Leverage (large size, high leverage)	0.650	0.492
12	Growth (small size, low leverage)	0.032	0.031
13	Growth (small size, high leverage)	0.060	0.054
14	Growth (large size, low leverage)	0.007	0.007
15	Growth (large size, high leverage)	0.017	0.016
16	Aggregate growth	0.019	0.019

Notes: Large/small firms are split at the median size. High/low leverage firms are split at the median leverage (debt-to-sales). Payout moments are computed for firms in the top decile by size.

4.3 Steady-State Equilibrium and Selection

Table 4 summarizes key equilibrium objects.

High-type firms own about 43 percent of active product lines and account for a similar share of GDP, yet they hold more than half of the debt stock and are substantially more levered than low types. Despite their higher leverage, high types pay lower average spreads because their continuation values are higher and they are less likely to default at a given debt level. Endogenous default reallocates roughly two percent of GDP per unit time, and the majority of this reallocation involves high-type firms. This pattern confirms that, in the calibrated economy, selection and reallocation are not purely about pruning low-productivity firms; financial distress also affects growth-oriented firms that are aggressive users of credit.

Figure 1 documents heterogeneity in innovation patterns. High types innovate about four times as intensively as low types of the same size. Within each type, innovation

Table 4. *Equilibrium Objects*

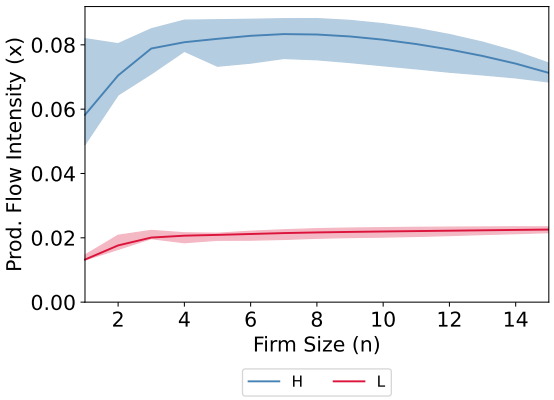
#	Statistics	Value
1	H-Type Share of GDP (%)	42.6
2	L-Type Share of GDP (%)	55.46
3	Inactive Share of GDP (%)	1.92
4	H-Type Endog. Default Realloc. (%)	1.38
5	L-Type Endog. Default Realloc. (%)	0.72
6	H-Type Share of Endog. Default Realloc. (%)	59.63
7	H-Type Share of Debt (%)	53.7
8	H-Type Share of Incumbents (%)	40.21
9	Agg. Growth (%)	1.9
10	Agg. Creative Destruction	0.0565
11	Incumbent Creative Destruction	0.0397
12	Entrant Creative Destruction	0.0168
13	Avg. (Incumbent) Innov. Intensity	0.0405
14	Avg. Innov. Intensity (H)	0.0696
15	Avg. Innov. Intensity (L)	0.0181
16	Avg. Debtor Leverage	38.07%
17	Avg. Debtor Leverage (H)	47.0%
18	Avg. Debtor Leverage (L)	31.2%
19	Avg. Spread	2.29%
20	Avg. Spread (H)	1.69%
21	Avg. Spread (L)	2.98%

intensity is roughly flat over size for sufficiently large firms, with some decline for the smallest ones due to the interaction of fixed costs and financial frictions. The right panel plots innovation intensity against leverage. Conditional on type, innovation declines as leverage increases: within-type, heavily indebted firms cut back on innovation despite tax advantages of debt. Together with the moments in Table 3, this implies that the empirical positive correlation between leverage and growth reflects selection: high-growth firms optimally choose higher leverage, but conditional on type, leverage is a drag on innovation.

Figure 2 shows the distribution of leverage and the drift of debt by size and type. Average leverage is relatively flat across size, with high types somewhat more levered than low types, particularly for intermediate sizes. The shaded regions reveal substantial heterogeneity within type-size cells, extending up to the debt limits. Despite high average leverage among high types, many high-type firms are actively deleveraging, while others lever up, depending on their distance to the debt limit and anticipated innovation oppor-

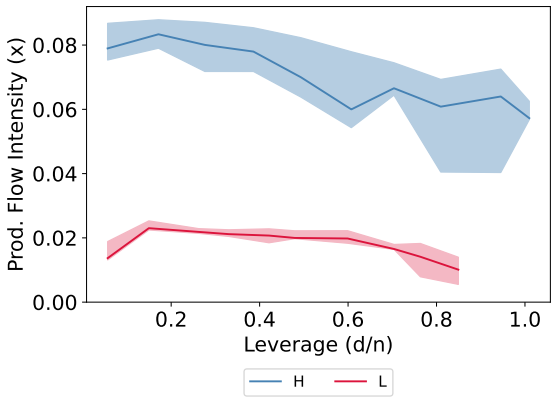
Figure 1. *Heterogeneity in Innovation*

(a) *Innovation across size*



Notes: Shaded regions represent the middle 90% of the distribution.

(b) *Innovation across leverage*

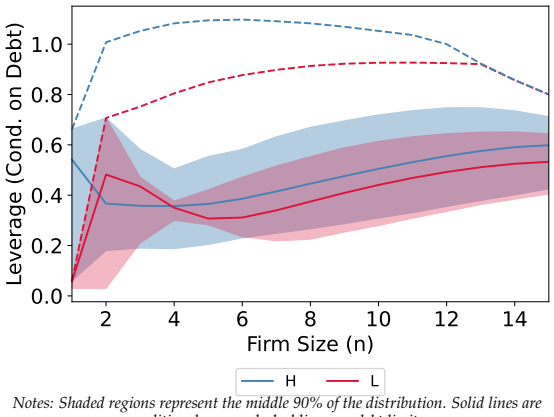


Notes: Shaded regions represent the middle 90% of the distribution.

tunities. Low-type firms tend to delever once they reach moderate size, reflecting lower innovation returns. The right panel shows that debt drift is positive for young, rapidly expanding firms but becomes negative for mature firms that have reached their optimal scale.

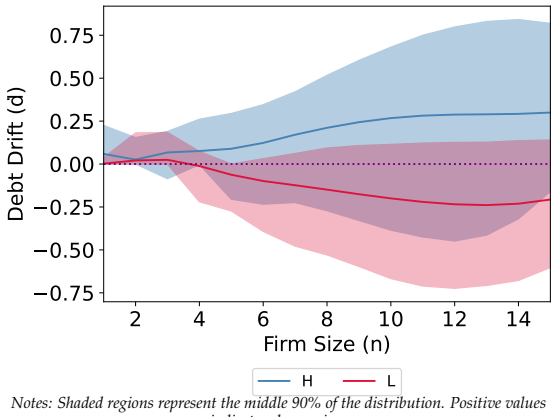
Figure 2. *Heterogeneity in Leverage*

(a) *Leverage across size*



Notes: Shaded regions represent the middle 90% of the distribution. Solid lines are conditional means; dashed lines are debt limits.

(b) *Debt drift across size*

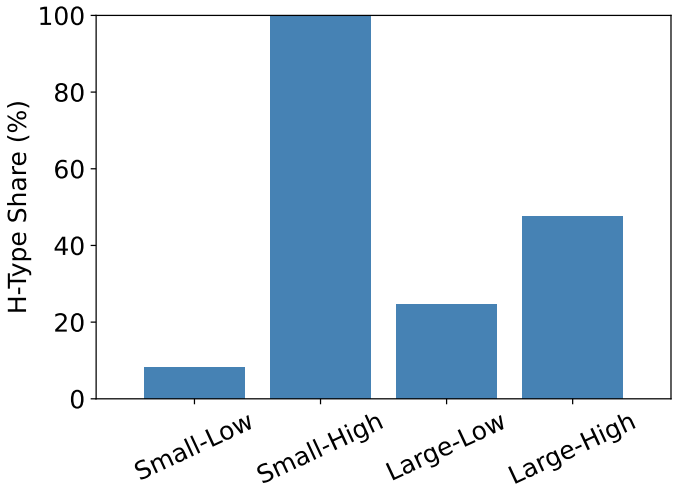


Notes: Shaded regions represent the middle 90% of the distribution. Positive values indicate releveraging.

Selection patterns across size and leverage are illustrated in Figure 3. High types are a minority among low-leverage firms, regardless of size. In contrast, they represent a majority among high-leverage firms, both for small and large firms. This pattern confirms the idea that high-growth firms endogenously select into higher leverage. It also suggests that financial shocks that disproportionately affect highly levered firms will hit high-type firms particularly hard and underscores the importance of understanding how

bankruptcy reallocates their assets.

Figure 3. *Selection across Size and Leverage*



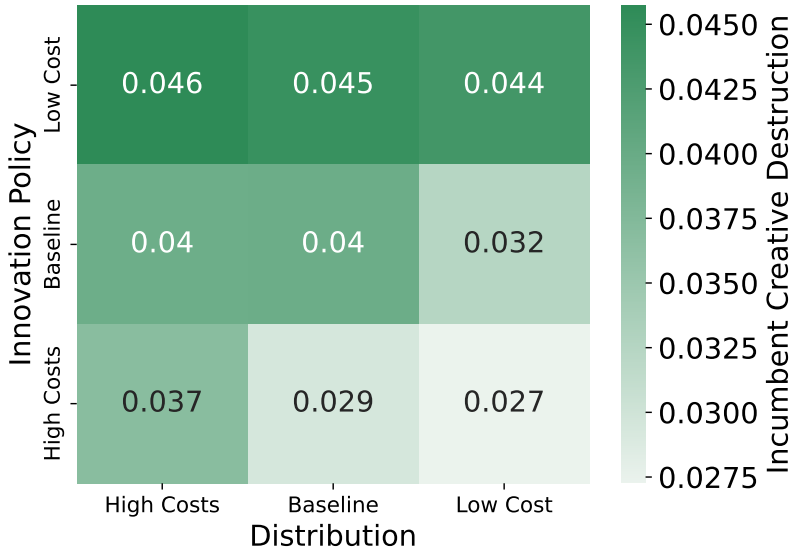
Notes: H-type share in four bins formed by size (small/large) and leverage (low/high).

4.4 Financial Frictions, Misallocation, and Creative Destruction

To understand the role of financial frictions for steady-state growth and misallocation, we consider counterfactual economies with higher or lower banking costs ν_0. Figure 4 plots incumbent creative destruction as a function of policies and firm distributions in three economies: baseline, high-cost, and low-cost. The diagonals report steady-state outcomes; off-diagonal entries combine the policy function of one economy with the firm distribution of another.

Increasing banking costs by 500 basis points reduces steady-state creative destruction from 0.040 to 0.037, a decline of about 8 percent relative to baseline. However, holding the distribution of firms fixed at its baseline values, the same increase in banking costs would reduce creative destruction to 0.029, a decline of nearly 28 percent. The difference between these two numbers reflects an endogenous compositional response: high-type firms reduce leverage and moderate their exposure to default in the high-cost steady state, partially insulating innovation from higher spreads. The cross-sectional distribution of balance sheets adjusts to dampen the impact of tighter financial conditions on growth. This adjustment is incomplete, leaving room for sizable real effects of financial frictions

Figure 4. *Borrowing Costs and Creative Destruction in Steady State*



Notes: Each cell shows incumbent creative destruction given a combination of firm policies (rows) and firm distributions (columns) across three steady states: baseline, high banking costs, and low banking costs.

on creative destruction and innovation, but it absorbs the majority of the policy effect.

The off-diagonal entries in Figure 4 reveal an asymmetry in how policy and distributional forces interact. The top row of the figure shows that when banking costs are low, creative destruction is nearly invariant to the distribution: incumbent innovation ranges only from 0.044 to 0.046 regardless of whether firms are drawn from the high-cost, baseline, or low-cost steady state. When credit is cheap, even firms with heavy debt loads can borrow at low spreads, so the distribution of balance sheets is largely irrelevant for innovation. The bottom row tells a different story. When banking costs are high, the distribution matters a great deal: creative destruction ranges from 0.037 (with the high-cost distribution, in which firms have already deleveraged) down to 0.027 (with the low-cost distribution, in which firms carry higher debt and are more vulnerable to spread increases). Adverse credit conditions thus interact with fragile balance sheets in a non-linear way, amplifying the damage to innovation when firms are caught with high leverage.

One cell deserves special attention: baseline policies combined with the high-cost distribution yield creative destruction of 0.040, the same as the baseline steady state. The

high-cost distribution, in which firms have deleveraged and reduced their exposure to default, is in some sense *better suited* for innovation when given access to normal credit terms. Firms that have paid down their debt can borrow at lower spreads and devote more resources to R&D. This observation foreshadows the crisis dynamics we study in Section 5: a crisis that forces surviving firms to delever can leave behind a distribution that is temporarily more favorable for innovation once credit conditions normalize, contributing to the productivity overshooting we document below.

The same experiments reveal that financial frictions induce misallocation in the sense that the dispersion in marginal returns to innovation across firms is larger when banking costs are higher. In the baseline economy, high-type firms innovate at nearly four times the intensity of low types (0.070 versus 0.018, Table 4), yet high types pay lower average spreads (1.69 percent versus 2.98 percent) because their higher continuation values reduce default risk. Within type, however, firms with weaker balance sheets face substantially higher spreads, driving a wedge between their shadow cost of funds and the risk-free rate. In the high-cost economy, this wedge widens: firms with strong innovation potential but high leverage operate with shadow costs of funds that are far above the risk-free rate, while some low-type firms with robust balance sheets face only modest spreads. From the perspective of a planner who could reallocate financial capacity across firms without affecting aggregate risk, this configuration would be dominated by one in which high-type firms face lower borrowing costs and invest more in R&D. The calibrated model thus implies that financial frictions not only reduce the average level of innovation but also distort its distribution across firms.

We can assess the magnitude of this misallocation by comparing the decentralized economy to a counterfactual in which borrowing costs are equalized across firms. In a simple experiment, we set $\delta(D, n, k)$ equal to its average across all debtors while keeping debt limits unchanged. This equalization raises creative destruction by roughly 15 percent and increases the contribution of high-type firms to aggregate growth. The gain is concentrated among the firms for whom financial frictions are most binding: high-type incumbents with above-median leverage, which in the baseline economy face spreads that are well above the type-average of 1.69 percent and innovate below their potential. Equalizing spreads effectively redirects credit capacity from low-type firms with ample debt headroom toward high-type firms that are constrained by steep spread schedules. The result confirms that the cross-sectional dispersion in borrowing costs, and the resulting misallocation of innovation effort, is a quantitatively important channel through

which financial frictions reduce aggregate growth.

5 Crisis and Transition Dynamics

We now use the calibrated model to study how a financial crisis affects growth, selection, and reallocation, and how the bankruptcy process shapes the recovery. We consider an unanticipated temporary increase in banking costs ν_0, which manifests itself as a spike in credit spreads. For computational tractability, the crisis takes the form of an MIT shock: it is unexpected, but once realized, its future path is deterministic.

Specifically, we increase ν_0 by 5 percentage points for three years, holding all other structural parameters fixed. This shock leads to higher spreads and tighter effective borrowing limits, especially for firms that are already close to their debt capacity. We compare the dynamic path of the economy under this shock to the baseline no-crisis path and decompose the response along several margins.

5.1 Aggregate Dynamics

Figure 5 shows the contemporaneous growth rate of GDP over time for the baseline and crisis paths. Growth falls immediately by about 15 percent relative to its steady-state level and remains depressed for roughly three years, while banking costs are elevated. During this period, firms face more expensive credit, and many reduce their innovation spending. The crisis also triggers a wave of defaults as firms with weak balance sheets and low innovation prospects can no longer service their debt.

Figure 6 plots the deviation of the log productivity index from its baseline path. On impact, the crisis drives productivity below baseline. Over the next several years, however, the path of productivity exhibits a mild overshooting: the economy converges back to the original growth path from above. This overshooting arises because the bankruptcy wave reallocates technologies toward stronger firms, and these firms subsequently invest more aggressively in innovation. The medium-run outcome of the crisis is thus a combination of a temporary level loss and a transient increase in the growth rate during the recovery phase.

Figure 5. *Contemporaneous GDP Growth*

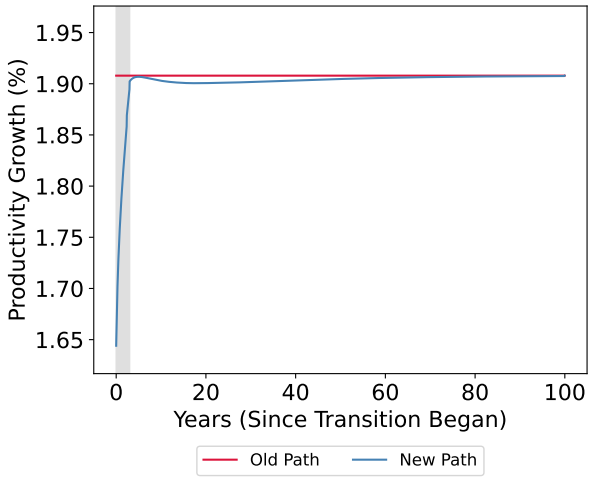
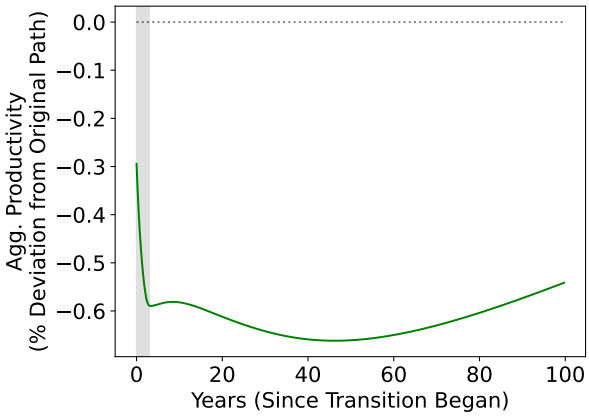


Figure 6. *Deviation of Productivity from Baseline*



5.2 Exit, Entry, and the Mass of Product Lines

The crisis affects not only aggregate growth but also the cross-sectional composition of firms. Figure 7 reports the mass of active and inactive product lines by size over time. On impact, the mass of active lines falls by about four percentage points relative to steady state. Because the bankruptcy resolution process is gradual, many lines remain temporarily inactive in bankruptcy for several years, leading to an elevated share of inactive lines.

At the same time, entry and post-crisis expansion gradually rebuild the mass of active lines, initially at small scales. The mass of small active firms rises above its steady-state level within a few years, whereas the mass of large active firms remains below baseline for more than a decade. The recovery is therefore driven by a reallocation from bankrupt and exiting incumbents toward a new cohort of growing firms. From the perspective of selection, the crisis accelerates the turnover of product lines, especially among highly levered low-type incumbents.

Figure 8 decomposes aggregate growth into the contributions of high-type incumbents, low-type incumbents, and entrants (which include firms that acquire product lines out of bankruptcy). The initial fall in growth is predominantly due to high types cutting their innovation. Because high types are both more levered and more reliant on external finance, the spike in spreads and tightening of debt limits has a disproportionately large effect on them. Low-type incumbents also reduce their modest innovation, but their contribution to aggregate growth is small to begin with.

Figure 7. *Mass of Product Lines*

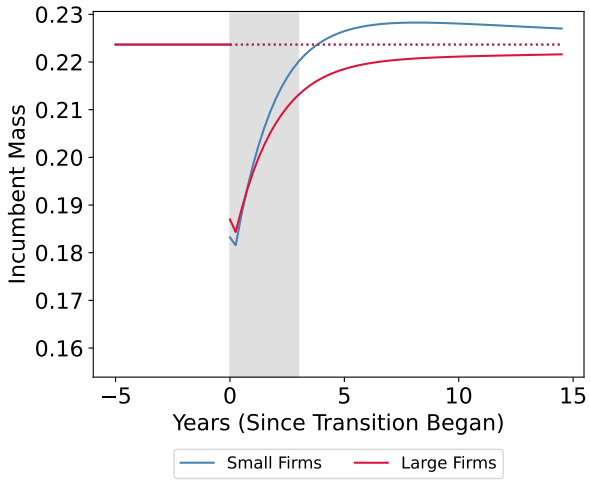
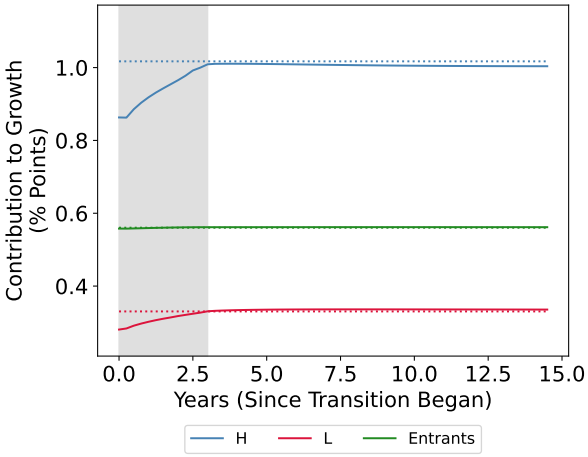


Figure 8. *Growth Decomposition*



As the economy moves into the recovery phase, the contribution of high-type survivors rebounds strongly. The composition of active firms shifts toward those that exited the crisis with healthier balance sheets and strong innovation potential, amplifying recovery. The entrant component, which captures innovation on product lines acquired through the bankruptcy process and through standard creative destruction, is positive and persistent. These lines tend to be held by firms with stronger balance sheets and higher innovation productivity, which amplifies their effect on growth relative to the lines they replace. Over the medium run, this component accounts for a sizable fraction of the cumulative deviation of productivity from baseline.

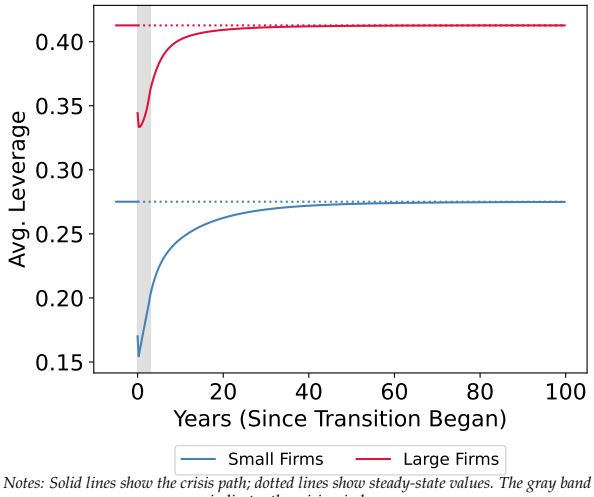
5.3 Cross-Sectional Dynamics of Leverage and Innovation

The crisis induces large changes in the cross-sectional distribution of leverage. Figure 9 plots the evolution of average leverage for small and large firms over the transition path, with dotted lines indicating steady-state values. Small firms delever dramatically: average leverage drops from about 0.28 to 0.15 on impact, a decline of nearly half, and takes more than 50 years to return to its pre-crisis level. The mechanism is twofold. First, the spike in spreads and tightening of debt limits force highly levered small firms into bankruptcy, removing them from the distribution. Second, surviving small firms actively pay down debt and reduce payouts, building a buffer against further credit shocks. Large firms experience a more moderate decline in leverage, from about 0.41 to 0.34, and recover within roughly 15 years. Their larger scale and lower fixed-cost burden provide a natural

buffer that limits the need for aggressive deleveraging.

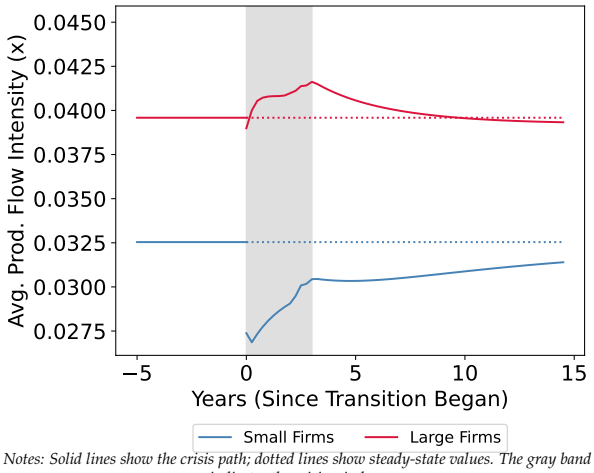
Figure 9. *Leverage Dynamics during Crisis*

(a) *Average leverage by size*



Notes: Solid lines show the crisis path; dotted lines show steady-state values. The gray band indicates the crisis window.

(b) *Average innovation intensity by size*



Notes: Solid lines show the crisis path; dotted lines show steady-state values. The gray band indicates the crisis window.

These balance-sheet adjustments interact with selection to shape innovation dynamics in a way that differs across the firm-size distribution. Figure 9(b) plots average innovation intensity for small and large firms. Small firms' average innovation drops sharply on impact, falling from about 0.033 to 0.027. This decline reflects both the intensive margin (individual firms tightening R&D budgets as spreads rise) and, more importantly, the extensive margin: highly levered, highly innovative small firms that default are replaced in the surviving distribution by a larger share of low-type incumbents that invest little in innovation.

The pattern for large firms is the opposite. Average innovation among large firms actually *rises* during and immediately after the crisis, peaking about 10 percent above its steady-state level. This occurs because the crisis disproportionately eliminates large firms with weak balance sheets and low innovation potential. The surviving pool of large firms is positively selected on both financial health and innovativeness, generating a temporary rise in average innovation intensity. This selection effect is the micro-level counterpart of the aggregate productivity overshooting documented in Figure 6.

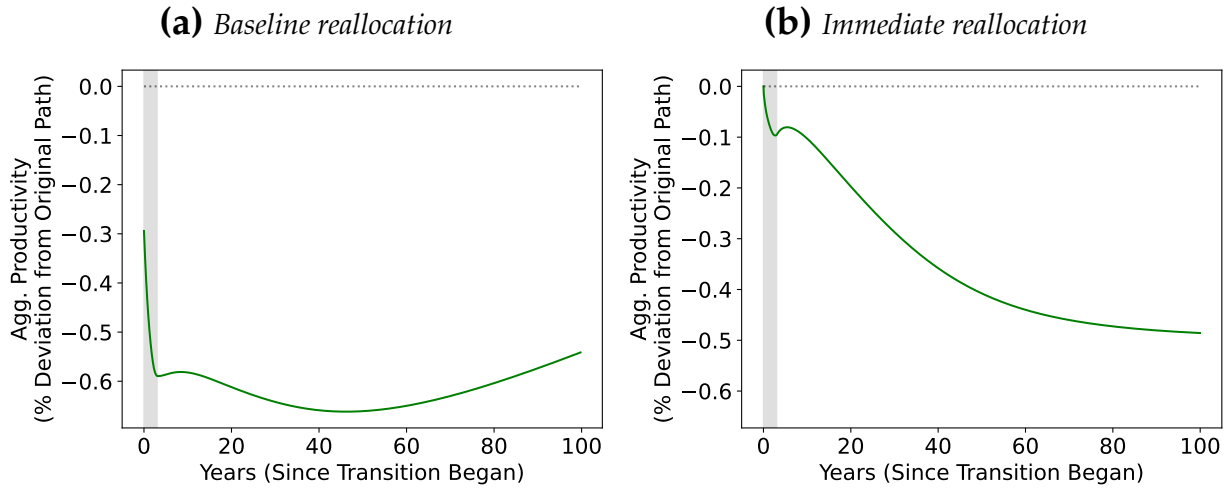
The compositional dynamics can also be seen by breaking out innovation by type within the small-firm category. High-type small firms that survive the crisis increase their innovation: with fewer competitors for new product lines, the expected return to innova-

tion rises, more than offsetting the higher cost of funds. Low-type small firms, by contrast, barely adjust their already low innovation effort. The decline in average small-firm innovation thus reflects a change in the composition of the small-firm pool, not a uniform cut in R&D across all surviving firms. This compositional channel has testable empirical implications: it predicts that average R&D intensity conditional on survival should rise among large firms after financial crises, even as aggregate R&D spending falls.

5.4 Role of Reallocation Speed

How much does the speed of the bankruptcy process matter? To answer this, we compare the baseline crisis to a counterfactual in which bankrupt product lines are reallocated instantaneously to new owners (keeping the distribution of buyers by type fixed). Figures 10a and 10b illustrate the difference.

Figure 10. *Growth and Reallocation Speed*



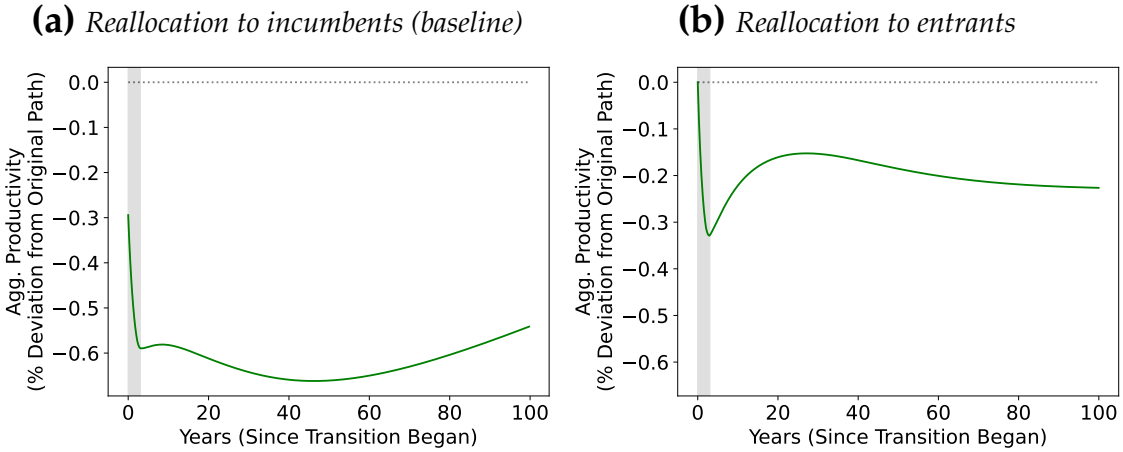
With immediate reallocation, the initial drop in productivity is smaller and the recovery is sharper. The intuition is straightforward. When product lines spend less time in bankruptcy, the economy wastes fewer resources in the form of idle technologies, and high-type buyers can start innovating on acquired lines sooner. The difference between the two paths provides a measure of the aggregate cost of legal and search frictions in bankruptcy. In our calibration, these frictions account for a substantial share of the level loss in productivity during the crisis, but they play a more modest role in shaping the long-run growth rate once the economy has returned to its balanced growth path.

5.5 Who Receives Bankrupt Assets? Incumbents vs Entrants

The impact of reallocation also depends on who acquires the bankrupt assets. In the baseline, product lines are sold to incumbents whose type is drawn from the cross-sectional distribution. We contrast this with a counterfactual in which bankrupt lines are always sold to entrants, who are more likely to be high type and start life with clean balance sheets.

Figures 11a and 11b compare the paths of productivity deviation under these two scenarios.

Figure 11. *Growth and the Type of Buyers*



When bankrupt assets are reallocated to entrants, the medium-run overshooting of productivity is much stronger, and the long-run deviation from baseline remains positive for a longer horizon. The reason is that entrants have a higher probability of being high type than incumbents in the cross-section, and they begin with no debt, giving them ample room to finance innovation on their new lines. In this sense, crises can be particularly beneficial for productivity when they are accompanied by stronger entry by innovative firms that can absorb distressed assets. This result echoes empirical findings that post-crisis recoveries are often led by young, high-growth firms rather than by pre-crisis incumbents.

6 Conclusion

This paper proposes a quantitative framework to study how firm-level financial frictions interact with endogenous innovation, selection, and crises. We build a continuous-time Schumpeterian growth model in which heterogeneous firms differ in size, innovation potential, and balance-sheet position, and in which banks endogenously set firm-specific interest rates and debt limits in response to default risk and costly bankruptcy. Calibrating the model to U.S. firm-level data on leverage, spreads, default, and growth, we find that financial frictions are quantitatively important for innovation and creative destruction, primarily by constraining high-growth, highly levered firms. Equalizing borrowing costs across firms raises creative destruction by about 15 percent in steady state, confirming that the cross-sectional dispersion in borrowing costs is quantitatively important for innovation and growth.

We then analyze the transition dynamics of a financial crisis with a temporary spike in intermediation costs. The crisis depresses innovation, triggers a wave of bankruptcies, and reduces growth in the short run. At the same time, the ensuing bankruptcy-driven reallocation can substantially strengthen the recovery: distressed product lines are transferred to firms with stronger balance sheets and higher innovation potential, which invest more aggressively in their expansion. Active reallocation through bankruptcy is thus a central driver of robust post-crisis growth in our model.

Our analysis further shows that the cleansing effect of crises is sensitive to institutional features of the bankruptcy process. Faster resolution and a greater likelihood that high-type firms acquire distressed assets both improve aggregate outcomes, highlighting the importance of efficient bankruptcy institutions for post-crisis recoveries.

From a broader perspective, our results suggest that understanding the joint distribution of firms' growth opportunities and balance sheets is central for thinking about crises and recovery. The positive correlation between leverage and growth in the data reflects selection: high-innovation firms borrow more because they have more to gain from expansion. This means that financial shocks disproportionately hit the firms that matter most for long-run growth. Whether the economy recovers strongly depends on how efficiently bankruptcy reallocates their assets toward firms with the capacity to innovate. Designing institutions that allow bankruptcy to serve as an efficient reallocation mechanism, rather than merely a costly legal process, is therefore central to harnessing creative

destruction in the face of financial turmoil.

A promising avenue for future work is to bring richer empirical evidence on bankruptcy resolution and post-bankruptcy firm performance to bear on the calibration of the reallocation process. Another is to use the framework to study stabilization policies—such as credit subsidies, loan guarantees, or macroprudential leverage caps—and to quantify the trade-off between short-run output stabilization and the long-run productivity costs of dampening the cleansing channel. Extending the model to incorporate heterogeneous lenders and macro-financial feedback would allow us to study how banking-sector fragility and regulatory capital requirements interact with firm-level innovation and selection. Finally, applying the model to specific historical crises, such as the Great Recession or the COVID-19 shock, and confronting its predictions with detailed micro data on distressed asset sales and post-crisis firm dynamics would further test the mechanisms emphasized here.

References

Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William R. Kerr, "Innovation, Reallocation, Growth," *The American Economic Review*, 2018, *108*, 3450–3491.

Achdou, Yves, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll, "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach," *National Bureau of Economic Research*, August 2017.

Aghion, Philippe and Gilles Saint-Paul, "Virtues of Bad Times: Interaction Between Productivity Growth and Economic Fluctuations," *Macroeconomic Dynamics*, 1998, *2*, 322–344.

_ , George-Marios Angeletos, and Abhijit Banerjee, "Volatility and growth: Credit constraints and the composition of investment," *Journal of Monetary Economics*, 2010, *57*, 246–265.

Akcigit, Ufuk and William R. Kerr, "Growth through Heterogeneous Innovations," *Journal of Political Economy*, August 2018, *126* (4), 1374–1443.

Barlevy, Gadi, "The Sullying Effect of Recessions," *Review of Economic Studies*, 2002, *69*, 65–96.

Blackburn, Keith, "Can Stabilisation Policy Reduce Long-Run Growth?," *The Economic Journal*, January 1999, *109*, 67–77.

Bornstein, Gideon, "A Continuous-Time Model of Sovereign Debt," *Journal of Economic Dynamics and Control*, September 2020, *118*, 1–26.

Bris, Arturo, Ivo Welch, and Ning Zhu, "The Costs of Bankruptcy: Chapter 7 Liquidation versus Chapter 11 Reorganization," *The Journal of Finance*, June 2006, *61* (3), 1253–1303.

Caballero, Ricardo J. and Mohamad L. Hammour, "The Cleansing Effect of Recessions," *American Economic Review*, December 1994, *84* (5), 1350–1368.

_ and _, "The Cost of Recessions Revisited: A Reverse-Liquidationist View," *Review of Economic Studies*, 2005, *72*, 313–341.

Cerra, Valerie and Sweta Chaman Saxena, "Growth Dynamics: The Myth of Economic Recovery," *American Economic Review*, March 2008, *98* (1), 439–457.

Chatterjee, Satyajit and Burcu Eyigungor, "The Firm Size-Leverage Relationship and Its Implications for Entry and Business Concentration," *Federal Reserve Bank of Philadelphia*, October 2020, pp. 1–33.

Cooley, Thomas F. and Vincenzo Quadrini, "Financial Markets and Firm Dynamics," *The American Economic Review*, December 2001, *91* (5), 1286–1310.

Cooper, Russell and John Haltiwanger, "The Aggregate Implications of Machine Replacement: Theory and Evidence," *American Economic Review*, June 1993, *83* (3), 360–382.

Foster, Lucia, Cheryl Grim, and John Haltiwanger, "Reallocation in the Great Recession: Cleansing or Not?," *Journal of Labor Economics*, 2016, *34* (1), S293–S328.

Geelen, Thomas, Jakub Hajda, and Erwan Morellec, "Debt, Innovation, and Growth," *SSRN*, July 2020.

Gopinath, Gita, Sebnem Kalemli-Ozcan, Loukas Karabarbounis, and Carolina Villegas-Sanchez, "Capital Allocation and Productivity in South Europe," *The Quarterly Journal of Economics*, 2017, pp. 1915–1967.

Jermann, Urban and Vincenzo Quadrini, "Macroeconomic Effects of Financial Shocks," *The American Economic Review*, February 2012, *102* (1), 238–271.

Klette, Tor Jakob and Samuel Kortum, "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, October 2004, *112* (5), 986–1018.

Laeven, Luc, Ross Levine, and Stelios Michalopoulos, "Financial innovation and endogenous growth," *J. Finan. Intermediation*, 2015, *24*, 1–24.

Li, Shaojin, Toni M. Whited, and Yufeng Wu, "Collateral, Taxes, and Leverage," *The Review of Financial Studies*, June 2016, *29* (6), 1453–1500.

Malamud, Semyon and Francesca Zucchi, "Liquidity, innovation, and endogenous growth," *Journal of Financial Economics*, 2019, pp. 519–541.

Miao, Jianjun, "Optimal Capital Structure and Industry Dynamics," *The Journal of Finance*, December 2005, *LX* (6), 2621–2659.

Osotimehin, Sophie and Francesco Pappada, "Credit Frictions And The Cleansing Effect Of Recessions," *The Economic Journal*, June 2017, *127*, 1153–1187.

Saint-Paul, Gilles, "Productivity Growth and the Structure of the Business Cycle," *European Economic Review*, 1993, *37*, 861–890.

A Derivations

A.1 Labor Market

Total labor demand is

$$\begin{aligned} L_t^D &= \int_{i\in\Phi} l_i di, \\ &= \int_{i\in\Phi} \frac{Y}{\lambda W} di, \\ &= \frac{||\Phi||}{\omega\lambda}, \end{aligned}$$

where $\omega = W/Y$. With inelastic labor supply L^S, labor market clearing $L^S = L^D$ implies

$$L^S = \frac{||\Phi||}{\omega\lambda}. \tag{15}$$

A.2 Final Good Production

Using $y_i = q_i Y/(\lambda W)$ and $W/Y = \omega$, final output can be written as

$$\begin{aligned} \ln Y &= \int_{i\in\Phi} \ln y_i di \\ &= \int_{i\in\Phi} \ln\left(\frac{q_i Y}{\lambda W}\right) di \\ &= \int_{i\in\Phi} \ln\left(\frac{q_i}{\lambda\omega}\right) di \\ &= \int_{i\in\Phi} \ln\left(q_i\right) di + \int_{i\in\Phi} \ln\left(\frac{1}{\lambda\omega}\right) di \\ &= \int_{i\in\Phi} \ln\left(q_i\right) di + \ln\left(\frac{1}{\lambda\omega}\right)\Phi \\ &= \int_{i\in\Phi} \ln\left(q_i\right) di + \ln\left(\frac{L^S}{\Phi}\right)\Phi. \end{aligned} \tag{16}$$

Let $\ln Q = \int_\Phi \ln(q_i)\, di$. Then

$$Y = Q\left(\frac{L^S}{\Phi}\right)^\Phi. \tag{17}$$

A.3 Normalized Firm Value

Normalizing by Y and using $\dot{Y}/Y = g_Y$, the normalized value $v = V/Y$ satisfies

$$\begin{aligned}(r_f - g_Y)\, v(d,n,k) \;=\; \max\Bigg\{0,\; \max_{x,p}\Bigg\{ & p\, n + \frac{\partial v}{\partial d}(d,n,k)\, \dot{d}(x,p;d,n,k) \\ & + x\, n\, [v(d,n+1,k) - v(d,n,k)] \\ & + \tau\, n\, [v(d,n-1,k) - v(d,n,k)] \\ & + 1_{k=H}\, \phi_{H\to L}\, [v(d,n,L) - v(d,n,H)] \\ & - \varepsilon\, v(d,n,k) \Bigg\}\Bigg\}.\end{aligned}$$

where the normalized debt drift is

$$\dot{d} = -\left[\frac{\pi n}{Y} - \chi - \frac{c(x)}{Y} - r(D,n,k)\, d\right](1-\tau_{tax}) + \frac{\varphi(p)}{Y} - g_Y d. \tag{18}$$

A.4 Firm Optimal Policies

Let ψ_D denote the Lagrange multiplier on the upper debt constraint. The first-order condition for payouts yields

$$0 = n + \left[\frac{\partial v}{\partial d}(d,n,k) + \psi_D\right][1 + 2\kappa p]\, n, \tag{19}$$

so

$$p^* = -\frac{1}{2\kappa}\left[1 + \frac{1}{\frac{\partial v}{\partial d}(d,n,k) + \psi_D}\right]. \tag{20}$$

The first-order condition for innovation intensity x is

$$0 = \left[\frac{\partial v}{\partial d}(d,n,k) + \psi_D\right]\eta\frac{1}{\theta_k}x^{\eta-1}n^{\eta(1-\sigma)} + n\,[v(d,n+1,k) - v(d,n,k)], \tag{21}$$

which implies

$$x^* = \left(-n^{1-\eta(1-\sigma)}\frac{\theta_k}{\eta}\frac{v(d,n+1,k) - v(d,n,k)}{\frac{\partial v}{\partial d}(d,n,k) + \psi_D}\right)^{\frac{1}{\eta-1}}. \tag{22}$$

B Computational Strategy

B.1 Steady State

Given a parameter vector, we solve for the steady-state equilibrium using nested fixed-point iterations. The outer loop searches for the creative destruction rate τ consistent with firm innovation and entry. Conditional on a guess for τ, the middle loop solves for firm value functions and optimal policies for each type, exploiting the one-way nature of transitions from high to low types.

Within the middle loop, an inner loop iterates on default decisions $B(D, n, k)$, debt limits $\bar{D}(n, k)$, and spreads $\delta(D, n, k)$ implied by the banking problem. For given default and debt limits, we compute spreads from (8). The innermost loop then solves the Hamilton–Jacobi–Bellman equation for firms, using a semi-implicit finite difference method with an upwind scheme as in Achdou et al. (2017). We exploit the one-way nature of type transitions to reduce the dimensionality of the system.

Once the value functions and policies are computed, we derive the stationary distribution of firms by solving the Kolmogorov forward equation, again using finite differences. Aggregating over the distribution delivers an implied rate of creative destruction. We update the outer-loop guess for τ until consistency is achieved.

B.2 Transition Dynamics

For the crisis experiments, we compute deterministic transition paths starting from the steady state. We guess a time path for creative destruction and for the wedge between GDP and consumption growth (which affects normalization). Given these guesses, we compute value functions and policy functions at each date using a fully implicit method. The fully implicit scheme is crucial for accurately capturing the evolution of debt limits and default decisions over time. Semi-implicit or explicit schemes tend to misalign default thresholds across dates and create spurious jumps.

Implementing the fully implicit method requires a Newton–Raphson algorithm for the value function, because today's optimal innovation and payout policies depend on today's value function. To avoid computing the full Jacobian, we use a sparse approxi-

mation that captures the main dependencies across neighboring grid points. Once policies are obtained at each date, we solve the forward Kolmogorov equation to obtain the evolving distribution of firms and recompute the implied creative destruction and growth paths. We iterate on the guessed paths until convergence.