Financing Business Dynamism: Leverage, Default, and Reallocation

Ufuk Akcigit Harun Alp Ege Y. Ercan*

September 1, 2026

Abstract

Growth-oriented firms borrow more, innovate more, and default more. We build a Schumpeterian growth model in which firms differ in size, leverage, and growth potential, and banks set firm-specific interest rates and debt limits, calibrated to U.S. supervisory loan-level data and Compustat. High growth-potential firms innovate roughly four times as intensively as low-potential firms and borrow more to finance this investment. Within type, more debt lowers innovation. The positive leverage–growth correlation therefore reflects selection, not a causal effect of debt. A temporary credit-cost shock triggers a default wave concentrated among levered high-potential firms and leaves productivity persistently below its no-shock path, and bankruptcy reallocation and survivor selection drive only a partial recovery. Temporary interest subsidies cushion the short-run loss, but the marginal firms they save are low-potential: a moderate subsidy improves the productivity path at every horizon, whereas an aggressive one, by stopping the default wave that fuels reallocation, eventually leaves the economy worse off than no intervention.

JEL Codes: G3, O4

Keywords: Productivity, Credit Shocks, Transition Dynamics, Debt, Financial Frictions, Default, Creative Destruction, Stabilization Policy

*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

In the spring of 2020, as credit markets seized and firms faced a sudden spike in the cost of borrowing, governments across the advanced economies moved quickly to keep them alive, deploying loan guarantees, interest subsidies, and forbearance on a scale rarely seen outside wartime. The instinct is a natural one: a credit crunch forces otherwise viable firms into bankruptcy, and heading off those failures seems self-evidently worthwhile. A generation earlier, however, Japan learned the other side of that instinct, when banks that kept credit flowing to distressed "zombie" borrowers slowed the reallocation of resources toward more productive firms and the economy stagnated for the better part of a decade (Caballero et al., 2008). Which force dominates when a credit shock strikes an economy whose growth is driven by innovation: the cleansing that shifts resources toward the firms with the best ideas, or the destruction of the very firms that carry them? Answering it requires knowing which firms borrow, why they borrow, and what innovation is lost when they default, in normal times and in crises alike. This paper answers it with a quantitative growth model disciplined by confidential firm-level credit data for the United States.

We embed 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 three questions. First, how distortive are firm-specific borrowing costs and debt limits for innovation and creative destruction along the balanced growth path (BGP)? Second, when credit abruptly tightens, how costly is the resulting wave of default and restructuring for the economy's productivity path, and through which margins (exit, bankruptcy reallocation, and survivors' innovation) does the economy recover? Third, can temporary credit subsidies during the shock improve that path, or do they blunt the very reallocation through which the recovery operates?

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 innovation type that captures 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.

A representative bank 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 micro data on U.S. non-financial firms. We combine supervisory loan-level data, which provide information on bank–firm exposures and interest-rate spreads, with Compustat data on public firms' payouts. We construct moments describing the joint distribution of leverage, spreads, and growth across size and leverage groups, as well as the employment share of entrants and aggregate productivity growth, and choose parameters to match these moments. We place particular emphasis on reproducing three empirical regularities: growth and leverage are positively correlated in the cross-section, interest-rate spreads rise with leverage, and small firms pay higher spreads.

The calibrated model displays three central features on the balanced growth path. First, selection, not debt, drives the leverage–growth correlation. High growth-potential firms innovate several times as intensively as low-potential firms in every size and leverage quartile, roughly four times on average, and borrow more to finance this investment, accounting for a disproportionate share of both aggregate growth and aggregate debt; yet, despite carrying more debt, they pay lower average spreads than low-potential firms (1.70 versus 2.98 percent) because their higher continuation values make them less likely to default at any given debt level. The positive leverage–growth correlation in the data reflects selection: the same growth potential drives both innovation and borrowing. Holding innovation type fixed, more debt actually reduces innovation, because high spreads and a tighter default region raise the shadow cost of internal funds. Second, firm-specific borrowing costs generate misallocation of innovation across firms. Third, bankruptcy and

reallocation play a modest but non-negligible role on the BGP. Endogenous defaults reallocate about one percent of product lines with debt per year, and roughly 60 percent of these reallocations involve high-potential firms, highlighting that financial distress affects precisely the firms that are central for long-run growth.

We then use the model to analyze a credit-cost shock: an unanticipated, temporary increase in intermediation costs that raises credit spreads, hitting hardest the firms most dependent on debt. The shock triggers an immediate fall in innovation and a wave of defaults concentrated among levered high growth-potential firms: cumulated over the first five years, excess defaults hit high-potential types roughly three-to-one relative to low-potential types. The shock leaves a persistent productivity-level loss: aggregate productivity is about 0.33 percent below its no-shock path when the shock ends and about 0.46 percent below at year fifteen, with the gap still widening at the end of the horizon; at no point does productivity rise above its no-shock path.

What recovery there is runs through bankruptcy reallocation and survivor selection. Product lines idled in bankruptcy, roughly 7.5 times their no-shock share at impact, are re-absorbed within about eight years; small firms lose about 24 percent of their product lines at impact, but their line mass is back on path by year four; and the firms left standing are the healthier ones, so measured innovation among survivors runs at or above the no-shock path (large firms up to 3 percent above, surviving small high-potential firms about 10 percent above in the first year). The scar is compositional: the growth contribution of small high-potential firms falls by more than half at impact and remains below its no-shock level at year fifteen, the micro counterpart of the widening aggregate gap.

The final part of the paper asks whether stabilization policy can improve on this outcome. We study a temporary interest subsidy available to all borrowers during the shock window. The subsidy reduces each borrower's net-of-tax interest expense by either 5 or 15 percentage points of its pre-tax interest expenses, financed by lump-sum taxes and compared with no intervention. Four results emerge. First, inside the shock window subsidies work as intended: at year three the productivity gap is -0.21 percent under the 15pp subsidy, -0.28 percent under the 5pp subsidy, and -0.33 percent with no intervention. Second, the defaults that subsidies prevent are overwhelmingly low-potential firms: even the aggressive 15pp subsidy barely dents high-potential defaults (cumulative excess defaults of 0.060 with no intervention versus 0.051, in units of product-line mass), while essentially eliminating excess low-potential defaults. Within five years, the share of low-potential lines that default falls from about 13 percent to 3 percent, against 21 to 19 percent

for high-potential lines. Third, saving low-potential firms congests reallocation: the 15pp subsidy cushions the impact drop in high-potential line mass (-13.4 versus -15.8 percent with no intervention) yet recovers slowest, leaving the lowest high-potential mass even fifty years out. Fourth, the ranking reverses within a decade: the 15pp path crosses the 5pp path around year 6.5 and the no-intervention path around year 10, and by year 25 it is the worst outcome (-0.72 percent, versus -0.58 with no intervention and -0.48 under the 5pp subsidy). Our results, therefore, suggest a tradeoff between the short-term stabilization and long-term growth. Moderate support dominates laissez-faire, while aggressive support, tempering the default wave and thus reallocation, eventually leaves the economy worse off than doing nothing.

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). This cleansing view has long been contested (Long, 1990). 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, that neither pole fully obtains. A credit-cost shock leaves a persistent productivity-level loss, yet the partial recovery that does occur runs through bankruptcy-driven reallocation and survivor selection. Reallocation limits the scarring without ever making the shock pro-growth.

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 the post-shock productivity path into the contributions of exiting firms, surviving incumbents of different sizes and growth types, and entrants acquiring reallocated product lines. Moreover, because recoveries, resolution delays, and the reallocation of bankrupt technologies are modeled explicitly, the framework provides a laboratory for stabilization policy: we use it to show that the ranking of temporary interest-rate subsidies depends critically on the horizon over which they are evaluated.

The paper proceeds as follows. Section 2 presents the model. Section 3 describes the data and calibration strategy and discusses identification. Section 4 analyzes the BGP and the role of financial frictions for innovation, selection, and misallocation. Section 5 studies the transition dynamics of a credit-cost shock (its aggregate productivity cost and the margins of default, reallocation, and survivor innovation behind it), quantifies how the speed of bankruptcy resolution shapes the recovery, and evaluates temporary credit-spread subsidies during the shock window. 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 in which 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

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

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 $T(t)$ denotes taxes and transfers. The usual no-Ponzi condition requires

$$\liminf_{t\to\infty} \exp\left(-\int_0^t r_f(s)\, ds\right) \left[A(t) + D(t)\right] \geq 0.$$

Household's maximization implies the standard Euler equation, 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 intermediate good product lines, whose total measure 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: the firm does not target specific varieties but improves a random one 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 turnover of producers is *creative destruction*, arriving at rate $\tau \geq 0$, and it 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 a productivity enhancement *intensity* of $x \geq 0$ (where $x = \frac{X}{n}$) is

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

where we take $\sigma = 1 - 1/\eta$, yielding constant returns to scale.

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 adjustment 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. The firm's cash flow from production, fixed costs, productivity investment, and payouts is then

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

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. However, if the firm is at its limit, then its debt drift must be non-positive, $0 \geq \dot{D}(x, p; \bar{D}(n, k), n, k)$.

At any moment, the firm can exit the economy. Exiting shareholders forgo the earnings and debt obligations of their firms. So, they choose to liquidate such that the value they get from the firm is never negative. In addition, the firm is forced into early liquidation whenever shareholder value falls below a threshold of $o_{close} \geq 0$ per product line. The debt limit is defined in equation (9) below.

Therefore, exit occurs as

$$B(D, n, k) = 1\{D > \bar{D}(n, k)\}, \tag{6}$$

that is, the firm defaults as soon as continuing is worth less than exiting (or the liquidation rule binds). A firm that loses its last product line ceases to exist and its shareholders receive nothing: $V(D, 0, k) = 0$. Because a firm with no lines cannot service any debt, $\bar{D}(0, k) = 0$, so $B(D, 0, k) = 1$ for any $D > 0$. When o_{close} is positive, default is triggered while continuation value is still positive. Solving the exit decision amounts to finding the debt levels at which continuation value falls below $o_{close}\, n\, Y$.

There are three types of bankruptcies. First, 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 turns 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 $B(D, n, L) = 1$. Second, 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 consists of a representative bank that borrows from the household and lends to individual firms. The bank observes each debtor's state and offers a loan menu 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 other banking costs.

At high enough debt levels, a firm cannot commit to repaying its loan: a small firm handed a large loan would optimally default at once, and no interest rate could compensate the bank for the loss. The bank denies such loans through debt limits. Interest-rate spreads and borrowing constraints are thus both set endogenously 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, $\delta(D, n, k) \geq 0$.

The bank funds this loan by borrowing from the household at the risk-free rate, r_f. It borrows at the risk-free rate because lending to a continuum of firms perfectly diversifies credit risk: total gains and losses are certain, so the bank always honors its obligations to the household.

The debtor firm defaults with some flow rate. When a bankruptcy occurs, the firm has

[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).

some size-dependent liquidation value, potentially allowing the bank to recover some funds. Let $\mathbb{E}[\Phi_D(D,n,k)] \geq 0$ denote the loan's expected loss rate: the flow rate of default times the fraction of principal the bank fails to recover. The bank also expends ν_0 per unit of loan extended, capturing other banking costs. When granting a loan D to the firm characterized by (n,k), the bank's profit is

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

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:

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

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.

To deny loans that would end in immediate default, the bank lends only at debt levels at which the firm's continuation value remains above the liquidation threshold. The debt limit is therefore

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

so the bank never extends a loan that would immediately trip the default rule.

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 states (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 is paid to the bank. At Poisson rate $\psi > 0$, the product lines are sold to a buyer and the original firm ceases to exist. 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 whose type is drawn randomly from incumbent firms' type distribution. 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, both threat points are zero, and the bank's bargaining power is $(1-\xi) \in [0,1]$. The most the buyer can pay 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 continuing is worth no more than the liquidation threshold. 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

$$\bar{D}(n,k)(1-\xi).$$

The bank does not know the future buyer's type when it extends credit to the debtor firm. Let $m_{B,H}$ denote the probability that the buyer is a high type, equal to the share of high types among incumbent firms. The bank's expected recovery from selling a firm of size $n > 0$ 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 remains subject to creative destruction during bankruptcy. The firm's value to the bank therefore solves

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

with $R_L(0) = 0$. This market structure ensures that endogenous defaults do not reduce contemporaneous output. As an alternative to the recovery value in (10), the bank can claim an outside option $\mu_0 Y \geq 0$, capturing recovery value not otherwise modeled.

The bank prices its loan conscious of the different liquidation values under different bankruptcy scenarios. Let $1_\phi = 1_{k=H} B(D, n, L)$ indicate whether the firm characterized by (D, n, k) would declare bankruptcy upon transitioning to a low type. Let $1_\tau = B(D, n-1, k)$ indicate whether the firm would default upon losing one product line. The bank's expected loss rate is then

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

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. Instead, the value recovered above debt's face value dissipates.

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 to capture a product line; upon entry, a firm holds that one line, carries no debt ($D = 0$), and draws a 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 technology mirrors the incumbent's, with 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.

Entrants choose their innovation intensity z to maximize expected value:

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

A successful entrant captures a product line from an incumbent, so entry is the final source of the creative destruction that incumbents face.

2.8 Firm Value

Recursively, firm value is

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

subject to $D\leq \bar{D}(n,k)$, $\dot{D}$ as in (5), and $V(D,n,k)=0$ when $D>\bar{D}(n,k)$.

The maximization is over investment and payout while the firm operates. The terms on the right-hand side capture, in order: the flow payout to shareholders; the change in firm value as debt evolves; the gain from a successful innovation, arriving at rate xn; the loss from creative destruction, at rate τn; and, for high types, the loss from transitioning to the low type, at rate $\phi_{H\to L}$. Per (9), the firm exits once its debt can no longer be rolled over, that is, once shareholder value would fall below $o_{close}\,n\,Y$.

The firm's first-order conditions are in Appendix A.4. Normalizing firm value and debt by output, $v\equiv V/Y$ and $d\equiv D/Y$, the optimal innovation intensity when the firm's debt is below its debt limit is

$$x^*=\left(-n^{1-\eta(1-\sigma)}\frac{\theta_k}{(1-\tau_{tax})\eta}\frac{v(d,n+1,k)-v(d,n,k)}{\frac{\partial v}{\partial d}(d,n,k)}\right)^{\frac{1}{\eta-1}}. \tag{13}$$

The key intuition is that $\frac{\partial v}{\partial d}(d, n, k)$ influences the innovation policy in addition to the standard $v(d, n+1, k) - v(d, n, k)$. The larger the shadow value of internal funds (i.e., the larger in magnitude $\frac{\partial v}{\partial d}$), the less innovation the firm undertakes.

2.9 Aggregate Growth and Equilibrium

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

$$g = \ln(\lambda)\tau \tag{14}$$

Definition A stationary balanced-growth equilibrium is a collection

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

such that y_i and p_i maximize product-line profits as in (3) and (2); W clears the labor market (4), with $||\Phi|| = 1$, of which $\int n\, d\Gamma$ is held by incumbents and the remainder is in bankruptcy; V solves (12) with debt drift (5) subject to $D \leq \bar{D}(n, k)$, with optimal policies x, p and default rule B (6); z maximizes entrant value (11); δ sets the bank's expected profit (7) to zero as in (8), $\bar{D}$ is the debt limit (9), and R_L solves (10) given the incumbent high-type share $m_{B,H} = \int 1\{k = H\}\, d\Gamma \big/ \int d\Gamma$; Γ is the stationary distribution of firms over (d, n, k), where $d = D/Y$ is debt normalized by output, induced by the optimal policies together with creative destruction, type transitions, default, bankruptcy resolution, and entry, i.e., the time-invariant solution of the associated Kolmogorov forward equation; creative destruction is consistent, $\tau = \int x\, n\, d\Gamma + z$, with $g = \ln(\lambda)\tau$ (14); and $r_f = \rho + \vartheta g$ from the Euler equation (1), with goods-market clearing $Y = C + \int c(x)\, d\Gamma + c_E(z)$, where banking costs, equity adjustment costs, fixed operating costs, and all bankruptcy-related recoveries and surpluses are transfers settled lump-sum with the household through T.

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, and growth, as well as aggregate growth, while preserving a transparent mapping from parameters to

moments.

3.1 Data Sources and Variable Construction

All moments except four 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 loan-level information on outstanding balances, interest rates, contractual spreads, internal risk ratings, and realized default events, which we aggregate to firm-year observations. 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. 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.

3.2 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 $\tau_{tax} = 0.35$, following 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).

We calibrate internal parameters to minimize the distance between model moments model(i) and data moments data(i),

$$\sum_{i=1}^{16} \frac{w(i)|\text{model}(i) - \text{data}(i)|}{\frac{1}{2}|\text{model}(i)| + \frac{1}{2}|\text{data}(i)|}, \tag{15}$$

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, the employment share of entrants, and aggregate growth. The weighting scheme penalizes proportional deviations, so that moments that are close to zero do not receive lower weight.

We explore the parameter space in three stages. First, we generate candidate parameter 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. (2022), with modifications to handle endogenous debt limits and spreads. The computational strategy is summarized in Appendix B.

3.3 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 early liquidation threshold o_{close} determine the extent to which banks recover in bankruptcy and thus shape debt limits, particularly for large firms. A lower 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 the pace of entry. The innovation productivity of entrants θ_E shapes the entry margin and the share of small firms; we identify it from the employment share of entrants.

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	0.35
5	ψ	Bankruptcy resolution rate	0.46

3.4 Calibrated Parameters

Table 2 reports the internally calibrated parameters. We 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	Buyer's bargaining share	0.421
5	o_{close}	Early liquidation threshold, per product line (in units of Y)	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 (gross of fixed costs) in one instant, roughly 38 percent 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 79 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 break-even condition.

On the financial side, the bargaining parameter ξ implies that, on average, assets sold out of bankruptcy trade at a discount of around 42 percent relative to their implied debt capacity. The fixed recovery value μ_0 corresponds to roughly 0.62 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 72 percent of the model's average spread among 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 productive at innovation as low types. Entrant innovation productivity is comparable to that of high types, reflecting aggressive innovation at entry that is necessary to match the observed employment share of entrants. The step size $\lambda = 1.40$ implies that the profit rate from each product line is about 29 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, and an individual high-type firm has an expected high-type spell of about 20 years. This structure generates a non-degenerate mixture of types in the stationary distribution and allows the model to reproduce the observed differences in growth rates across size and leverage groups.

3.5 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.

Spreads are matched well across size and leverage groups, with one exception: for

Table 3. *Targeted Moments: Model vs Data*

#	Moment	Data	Model
1	Fraction of positive payouts (top size decile)	0.838	0.842
2	Mean payout conditional on positive	0.060	0.050
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	Employment share of entrants	0.020	0.017
8	Leverage (small size, low leverage)	0.062	0.080
9	Leverage (small size, high leverage)	0.616	0.610
10	Leverage (large size, low leverage)	0.070	0.139
11	Leverage (large size, high leverage)	0.650	0.513
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 (Compustat). The employment share of entrants is the employment share of age-zero firms from the Business Dynamics Statistics, averaged over 2013–2019.

large, highly levered firms, the model overpredicts the spread (3.0 percent versus 1.5 percent in the data) and underpredicts leverage (0.51 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 model also somewhat overpredicts leverage in the two low-leverage cells (0.080 versus 0.062 for small firms; 0.139 versus 0.070 for large firms). The employment share of entrants is matched closely (0.017 in the model versus 0.020 in the data). 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.

4 Financing and the Allocation of Innovation

In this section, we examine the main properties of the BGP equilibrium. We first document how innovation, leverage, spreads, and default vary across firms, and then quantify the role of borrowing costs for creative destruction by comparing BGP economies with higher and lower banking costs.

4.1 Equilibrium and Selection

Table 4 summarizes key equilibrium objects.

Table 4. *Equilibrium Objects*

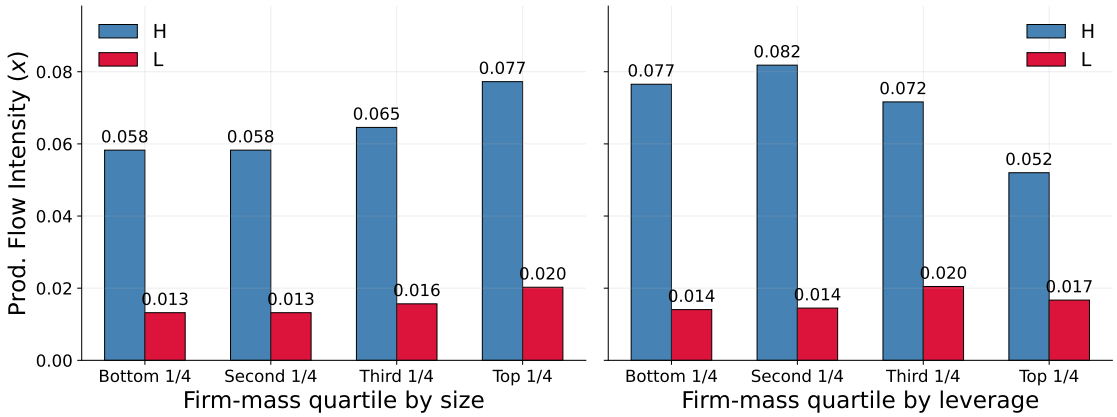
#	Statistics	Value
1	H-Type Share of GDP (%)	42.7
2	L-Type Share of GDP (%)	55.4
3	Share of Lines in Bankruptcy (%)	1.90
4	Endog. Default Realloc. (% of debtor lines)	1.01
5	H-Type Endog. Default Realloc. (% of H debtor lines)	1.39
6	L-Type Endog. Default Realloc. (% of L debtor lines)	0.72
7	H-Type Share of Endog. Default Realloc. (%)	59.9
8	H-Type Share of Debt (%)	53.1
9	H-Type Share of Incumbents (%)	40.3
10	Agg. Growth (%)	1.9
11	Agg. Creative Destruction	0.057
12	Incumbent Creative Destruction	0.040
13	Entrant Creative Destruction	0.017
14	Avg. (Incumbent) Innov. Intensity	0.040
15	Avg. Innov. Intensity (H)	0.070
16	Avg. Innov. Intensity (L)	0.018
17	Avg. Debtor Leverage (%)	38.1
18	Avg. Debtor Leverage (H, %)	46.6
19	Avg. Debtor Leverage (L, %)	31.6
20	Avg. Spread (%)	2.30
21	Avg. Spread (H, %)	1.70
22	Avg. Spread (L, %)	2.98

High types borrow the most yet pay the lowest spreads, a first sign of selection. They account for about 43 percent of GDP, yet 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 defaults reallocate about 1 percent of debtor product lines per year: H-type debtor lines at 1.39 percent and L-type debtor lines at 0.72 percent, so high types account for about 60 percent of the flow. About 1.9 percent of product lines are in bankruptcy at any time and thus do not contribute to R&D.

Figure 1 documents heterogeneity in innovation patterns. High types innovate several times as intensively as low types in every size and leverage quartile, roughly four times on average; in the bottom size quartile, the per-line innovation rate is about 0.058 for high types versus 0.013 for low types. Within high types, innovation rises with size, reflecting the interaction of fixed costs and financial frictions at small scale. The right panel plots innovation intensity across leverage quartiles. For high types, innovation declines as leverage rises beyond the second quartile, from about 0.082 to 0.052: heavily indebted high-type firms cut back on innovation despite the tax advantage of debt. Together with the moments in Table 3, this implies that the empirical positive correlation between leverage and growth reflects selection: high-potential firms optimally choose higher leverage, while for high types, the firms that do the innovating, high leverage is a drag on innovation.

Figure 1. *Heterogeneity in Innovation*



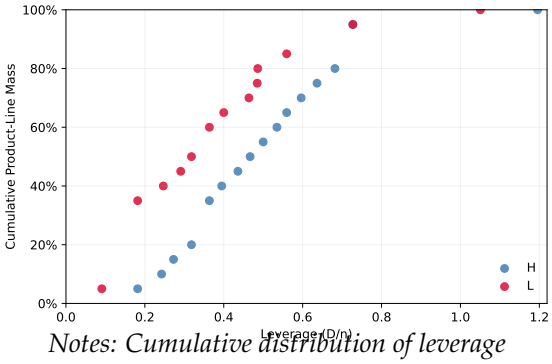
Notes: Per-line innovation intensity (arrival rate x, annual), by firm-mass quartile of size (left panel) and of leverage (right panel), for high (H) and low (L) growth-potential types. Boundary masses are split fractionally across quartiles.

Figure 2 documents the leverage distribution and the selection behind it. The left panel plots the cumulative distribution of leverage (debt-to-size, D/n) by type: the high-type CDF lies everywhere to the right of the low-type CDF, so at each quantile high types

carry more leverage. They are more levered throughout the distribution, not merely on average. The right panel shows the selection driving this gap: within each size class the high-type share rises steeply with leverage, from about 8 percent in the small, low-leverage bin to essentially 100 percent in the small, high-leverage bin, and from roughly a quarter to a half among large firms. Growth potential and leverage thus move together in the cross-section, so a shock that falls on highly levered firms falls disproportionately on high types, precisely the firms that matter most for growth.

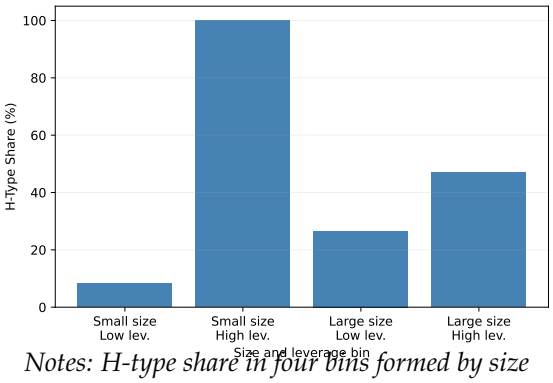
Figure 2. *Distributions*

(a) *Leverage CDF*



Notes: Cumulative distribution of leverage (debt-to-size, D/n) across product lines, plotted separately for high- (H) and low-type (L) firms.

(b) *Selection across Size and Leverage*



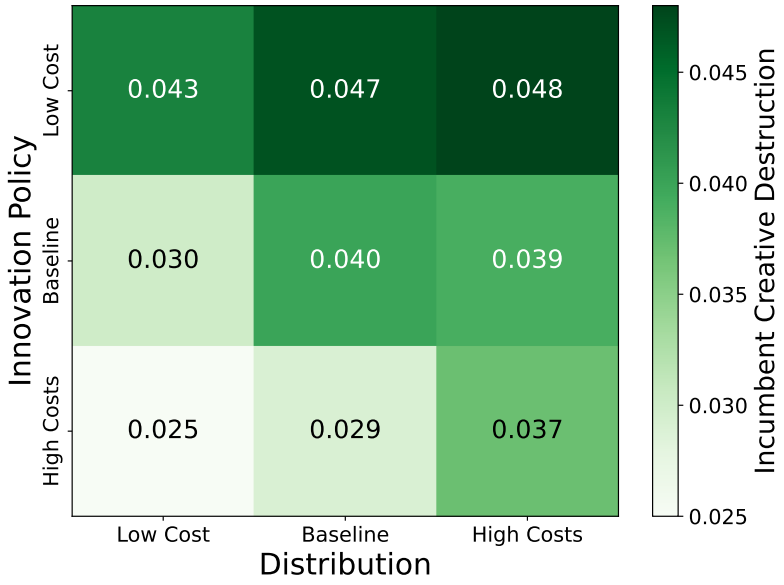
Notes: H-type share in four bins formed by size (small/large) and leverage (low/high).

4.2 Financial Frictions and Creative Destruction

To understand the role of financial frictions for growth and misallocation on the BGP, we consider counterfactual economies with higher or lower banking costs ν_0. Figure 3 plots incumbent creative destruction as a function of policies and firm distributions in three economies: baseline, high-cost, and low-cost. The diagonals report BGP 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 incumbent creative destruction on the BGP 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:

Figure 3. *Borrowing Costs and Creative Destruction on the Balanced Growth Path*



Notes: Rows: innovation policy functions from the low-cost, baseline, and high-cost BGPs; columns: firm distributions from the same three BGPs. Each cell reports the incumbent creative-destruction rate (annual); diagonal cells are the BGP outcomes.

high-type firms reduce leverage and moderate their exposure to default in the high-cost BGP, 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 sizeable real effects of financial frictions on creative destruction and innovation, but it absorbs the majority of the direct effect of higher costs.

The off-diagonal entries in Figure 3 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 about 0.043 to 0.048 regardless of whether firms are drawn from the high-cost, baseline, or low-cost BGP. 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 0.025 (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 0.039, essentially the same as the baseline BGP (0.040). The high-cost distribution, in which firms have deleveraged and reduced their exposure to default, performs on par with the baseline distribution once firms are given access to normal credit terms: deleveraged balance sheets are not themselves a long-run drag on innovation. This observation matters for the transition dynamics we study in Section 5. It implies that the persistent productivity loss we find there does not stem from the deleveraged balance sheets the shock leaves behind; it stems from the transition itself, in which the balance-sheet adjustment is forced to occur all at once, through a wave of bankruptcies that idles technologies and removes levered firms with high growth potential. The same comparison also anticipates why the recovery, though partial, runs through bankruptcy reallocation and survivor selection: once credit conditions normalize, surviving firms with strong balance sheets can borrow at low spreads and devote more resources to innovation.

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 (1.70 percent versus 2.98 percent) because their higher continuation values reduce default risk. Within type, however, the bank's pricing schedule charges firms with weaker balance sheets higher spreads, driving a wedge between their shadow cost of funds and the risk-free rate. In the high-cost economy this wedge widens, so that a firm with strong innovation potential but high leverage can face a higher cost of funds than a low-potential firm with a robust balance sheet. 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 innovation. The calibrated model thus implies that financial frictions not only reduce the average level of innovation but also distort its distribution across firms.

5 Transition Dynamics of a Credit-Cost Shock

We now use the calibrated model to study how a credit-cost shock 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 spikes credit spreads. For computational tractability, it takes the form of an MIT shock: it is unexpected, but once realized, its future path is deterministic. Spreads and debt limits continue to reprice endogenously along the entire transition path.

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, biting hardest for firms already close to their debt capacity. Throughout, we compare the economy's dynamic path to the *no-shock baseline path*, the path it would have followed absent the shock, and, unless noted otherwise, report transition objects as percent deviations from that path (figure axes labeled "No-Shock Path" or "Original Path" refer to this object; the gray band in each figure marks the three-year shock window). We first quantify the shock's aggregate productivity cost; we then trace the reallocation behind it: who defaults, how many product lines are idled in bankruptcy, and how quickly they are re-absorbed. Next, we document the cross-sectional dynamics of leverage and innovation and which groups account for the lost growth. Finally, we evaluate stabilization policy, comparing temporary credit-spread subsidies of different sizes against the shock path without intervention.

5.1 Aggregate Dynamics

Figure 4 plots the gap between aggregate productivity and the no-shock baseline path over the first fifteen years, in percent. On impact, credit becomes more expensive, debt limits tighten, firms cut innovation spending, and a default wave hits weak balance sheets. Aggregate growth falls sharply while banking costs are elevated and recovers most of the way once the shock ends at year 3. The productivity level, however, does not recover: the gap reaches about -0.33 percent when the shock ends and about -0.46 by year 15, still widening slowly at the end of the horizon.

The persistent loss operates through three channels: the default wave sends a large mass of product lines into bankruptcy, where they are idled (no innovation takes place

Figure 4. *Productivity Gap from the No-Shock Path*



Notes: Deviation of the aggregate (log) productivity level from the no-shock baseline path, in percent.

on a line awaiting resolution); costlier credit squeezes surviving incumbents' innovation spending precisely when internal funds are scarce; and the value of expanding into new product lines falls, discouraging innovation by incumbents and entrants alike. Equally important is what the figure does *not* show: at no point does productivity rise above the no-shock path. The recovery that does occur, in the *growth rate* rather than the level, runs through bankruptcy reallocating idled lines to new owners and through positive selection among survivors, margins the next two subsections document.

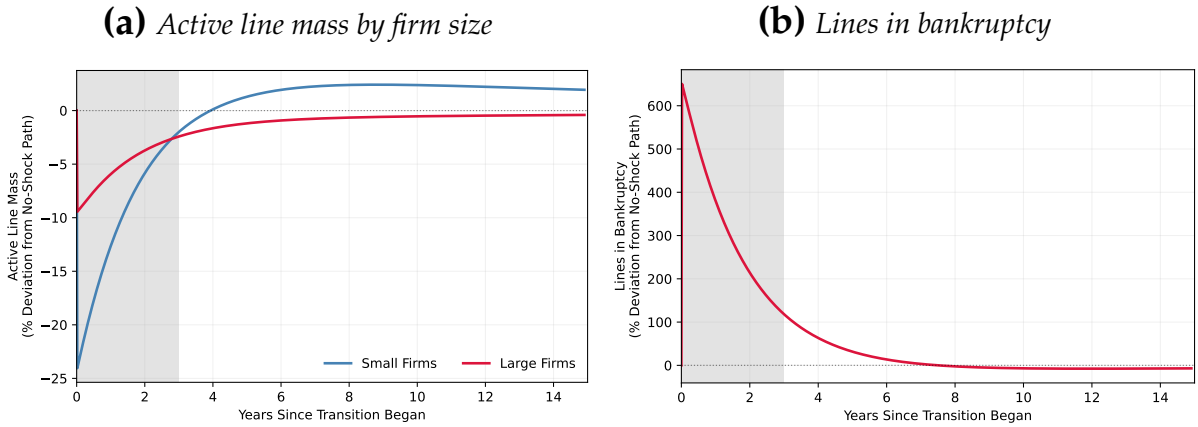
5.2 Exit, Entry, and the Mass of Product Lines

The aggregate loss conceals a highly uneven incidence. Figure 5(a) plots the mass of active product lines held by small and large firms. The shock is, first and foremost, a small-firm event: small firms lose about 24 percent of their lines at impact, against roughly 10 percent for large firms. Small-firm line mass is back on its no-shock path by about year 4 and temporarily rises 2–3 percent above it as vacated lines are re-absorbed, while large-firm mass converges back from below. The line-mass recovery is thus a reallocation of idled technologies toward small, growing owners.

The default wave's composition matters as much as its size. Exit concentrates among levered firms and, because high growth-potential firms optimally borrow more (Section

4.1), falls disproportionately on *high types*: cumulative excess defaults (defaults over and above the no-shock path) hit high types roughly three-to-one over the first five years (0.060 versus 0.021, in units of product-line mass), and about 21 percent of high-type lines default within five years, against about 13 percent of low-type lines (Figure 11 reports these numbers across policy scenarios; the values quoted here are those with no policy intervention). The shock thus destroys precisely the firms that matter most for growth, a point central to Section 5.5.

Figure 5. *Product Lines during the Transition*



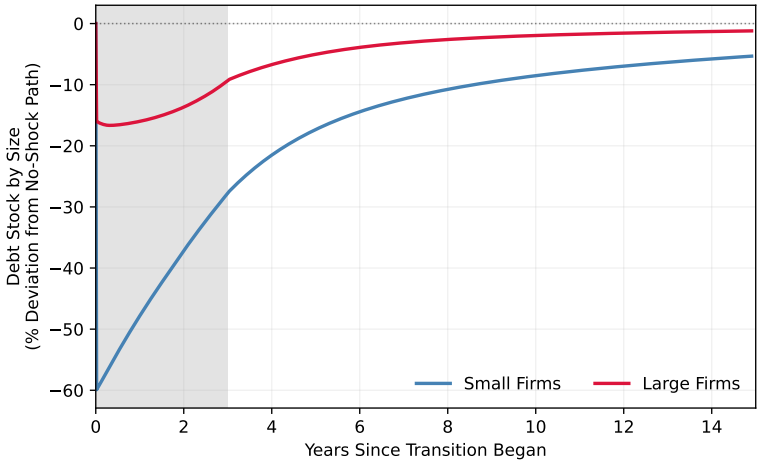
Notes: Panel (a): active product-line mass by firm size, percent deviation from the no-shock baseline path. Panel (b): share of lines idled in bankruptcy; a percent deviation of a small baseline share, not percentage points.

Because bankruptcy resolution is gradual, product lines are not instantly reallocated. Figure 5(b) shows the share of lines idled in bankruptcy jumping to roughly 7.5 times its no-shock level at impact and fully re-absorbed only by about year 8. While a line sits in bankruptcy it is not innovated on; idle technologies, together with scarce credit, are the shock's productivity bill. Credit itself contracts sharply: in Figure 6, the debt stock of small firms falls by about 60 percent at impact (large firms: about 17 percent); the small-firm stock is still about 27 percent below path when the shock ends and about 6 percent below at year 15. Bankruptcy does reallocate technologies, but slowly, and credit rebuilds slowly with it.

5.3 Cross-Sectional Dynamics of Leverage and Innovation

The shock also reshapes survivors' balance sheets. Figure 7(a) plots average leverage (debt to sales) by firm size. Small firms delever dramatically: their leverage ratio falls by

Figure 6. *Debt Stock during the Transition*



Notes: Aggregate debt stock of small and large firms, percent deviation from the no-shock baseline path.

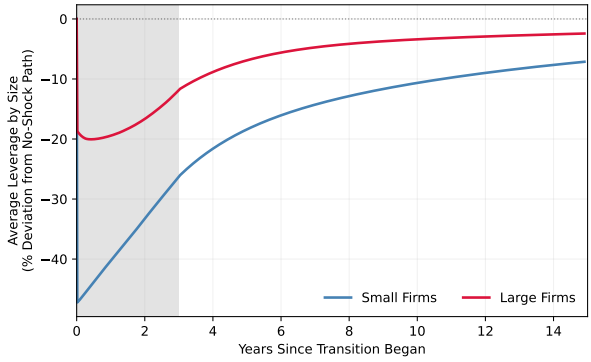
about 47 percent at impact, roughly halving, and is still about 7 percent below path at year 15. The mechanism is twofold: highly levered small firms default out of the surviving distribution, and survivors actively pay down debt and cut payouts to buffer against the elevated cost of credit. Large firms delever by roughly half as much (their ratio troughs about 20 percent below path) and are within a few percent of baseline by year 15. Small firms thus take both the larger line losses and twice the deleveraging, remaining cautious borrowers for a decade after the shock.

Balance-sheet adjustment interacts with selection to shape innovation differently across the size distribution. In Figure 7(b), average innovation among small firms falls by about 15 percent in the first year and is still about 4 percent below path at year 15. Among large firms, average innovation *rises* during and immediately after the shock, running up to about 3 percent above path: the surviving pool of large firms is positively selected on financial health and innovation potential.

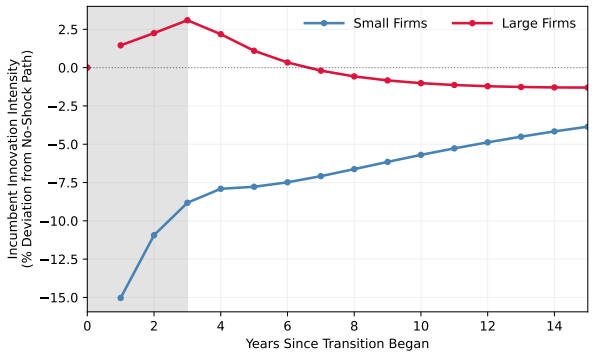
The fall in average small-firm innovation is not survivors cutting back; it is the mix of firms changing. Figure 7(c) splits small-firm innovation by type: conditional on type, surviving small firms innovate at or above path throughout, with high-type survivors about 10 percent above path in the first year while low types barely adjust their already low effort. What pulls the average down is that the high-type small firms, the ones that actually innovate, default the most and drop out, leaving a group tilted toward low types that barely innovate.

Figure 7. *Leverage and Innovation Dynamics during the Transition*

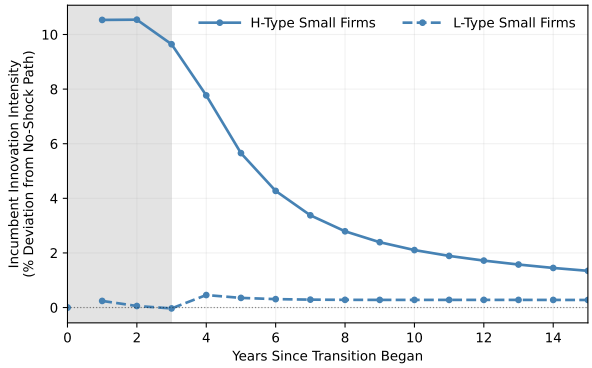
(a) *Average leverage by size*



(b) *Average innovation intensity by size*



(c) *Small-firm innovation by type*



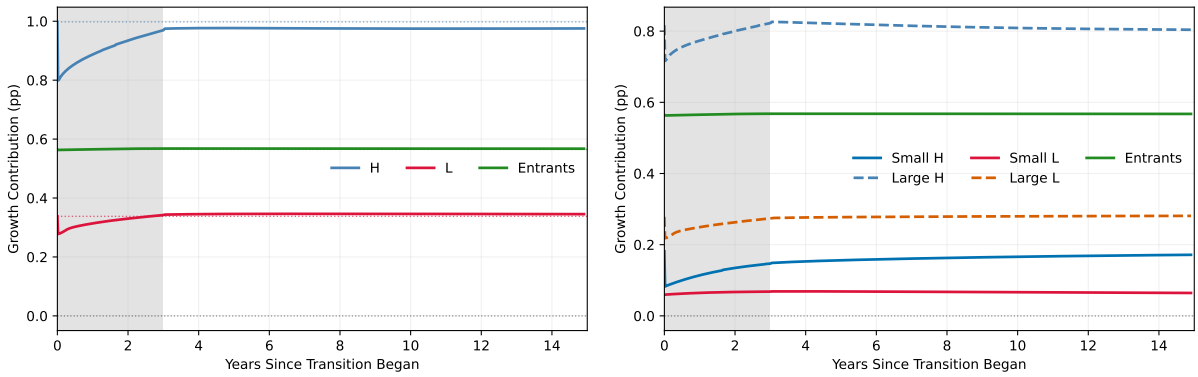
Notes: Percent deviation from the no-shock baseline path; for leverage, a percent change of the debt-to-sales ratio, not percentage points. Panel (c): surviving small firms' innovation by type; the aggregate small-firm decline in panel (b) reflects composition and mass loss, not survivors cutting back.

Which firms account for the lost growth? Figure 8 splits the aggregate growth rate into each group's contribution in percentage points, which sum to the 1.9 pp balanced-growth-path rate. Panel (a) shows that high types, who supply about half of that rate and are more levered and more reliant on external finance, drive almost the entire loss: their contribution falls at impact, pulling aggregate growth from about 1.9 to 1.65 pp, then recovers to just below its no-shock level by the end of the shock window. Low types dip and return to about their baseline, and entrants are essentially unchanged. Panel (b) locates the scar within the high types: small high-type firms' contribution falls by more than half at impact and remains below its no-shock level fifteen years out, the deepest and most persistent loss of any group. Large firms recover, and small low types edge up as they inherit vacated lines and face fewer rivals. Two caveats temper this low-type resilience: low types are not spared by the default wave (about 13 percent of their lines

default within five years, though levered high types suffer roughly three times the excess defaults), and their gains are shares of a *shrunken* pie, reflecting substitution along vacated lines rather than low types outgrowing high types. The missing contribution of small high-potential firms is the scar the shock leaves, the micro-level counterpart of the persistent productivity gap in Figure 4.

Figure 8. *Growth Decomposition during the Transition*

(a) *By type and entrants* **(b)** *Incumbents by size and type, and entrants*



Notes: Each group's contribution to aggregate growth, in percentage points; the groups sum to the aggregate rate of about 1.9 pp on the balanced growth path, and panel (a) marks each group's no-shock contribution with dotted lines.

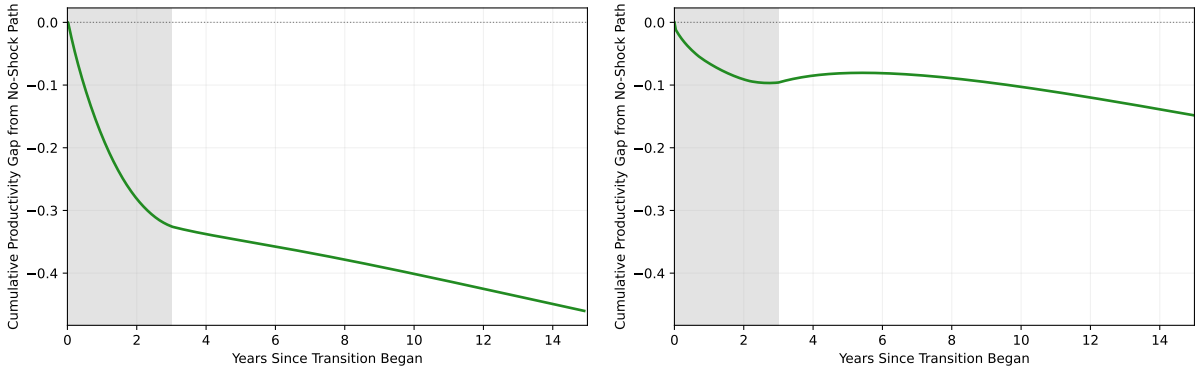
5.4 The Role of Reallocation Speed

How much of the productivity cost reflects slow resolution of defaulted product lines? To answer this, we compare the baseline transition to the same shock in an economy in which bankruptcy resolves almost immediately. In the calibrated economy, the resale/resolution rate is $\psi = 0.46$, so a defaulted line remains in bankruptcy for about $1/\psi \approx 2.2$ years on average. In the counterfactual economy, ψ is set essentially infinite, so bankrupt lines are reallocated almost immediately.

Figure 9 shows the gap between the shocked economies' aggregate productivity and their own counterfactual, unshocked paths. Immediate reallocation makes the loss far smaller: the productivity gap at the end of the shock window is about 0.10 percent instead of 0.33, and at year fifteen about 0.15 instead of 0.46. The mechanism is the one behind the aggregate productivity gap in Figure 4: when lines spend less time in bankruptcy, fewer technologies sit idle for innovation, and high-type buyers begin innovating on their acquired lines sooner. The gap between the two paths measures the aggregate cost of the

Figure 9. *Productivity Gap and Reallocation Speed*

(a) *Baseline resolution* **(b)** *Immediate reallocation*



Notes: Aggregate productivity, percent deviation from the no-shock baseline path. Panel (a): baseline bankruptcy resolution ($\psi = 0.46$, about two years in bankruptcy). Panel (b): bankrupt lines reallocated immediately.

legal and search frictions in bankruptcy. In our calibration these frictions account for a sizeable part of the productivity-level loss during the shock, but they matter far less for the long-run growth rate, which returns to its balanced-growth value once the idled lines are re-absorbed.

5.5 Stabilization Policy: Interest-Rate Subsidies

The default wave destroys high-type firms and idles their technologies, yet the same defaults reallocate product lines and set the recovery in motion, a sharp trade-off for stabilization policy. A natural policy, adopted in many crises, is to subsidize firms' interest expenses. We compare two temporary subsidies during the three-year shock window, each rebating a fixed fraction of the firm's interest expenses: 7.7 percent or 23.1 percent. These rates are equivalent to raising the debt tax shield by 5 and 15 percentage points, so the figures label the policies the "5pp" and "15pp" subsidies.[2]

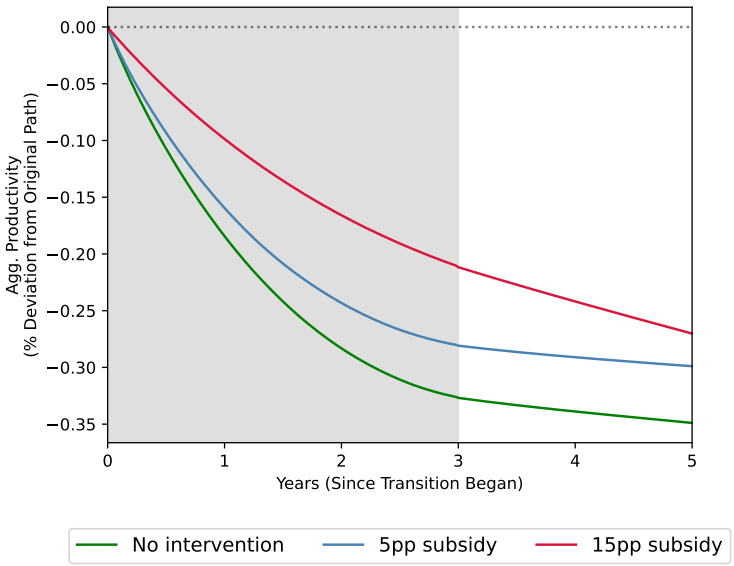
Firms with more debt and a higher default probability, and thus higher spreads, benefit more. Like the shock itself, the policy is unanticipated and operates only through the credit market: we abstract from any bank-capital channel, and the lump-sum financing enters the household budget constraint without distorting prices. We rank each subsidy

[2]With corporate tax rate τ_{tax}, a rebate at rate s on interest expenses lowers the after-tax interest cost from $(1 - \tau_{tax})\, rD$ to $(1 - \tau_{tax} - \Delta)\, rD$ with $\Delta = s\,(1 - \tau_{tax})$; the rates $s = 0.05/0.65 = 7.7$ percent and $s = 0.15/0.65 = 23.1$ percent thus correspond to Δ of 5 and 15 percentage points.

against the no-intervention path above, using the same productivity gap as in Figure 4.

Short-run stabilization. Inside the shock window, the subsidies work exactly as intended: they soften the rise in spreads and relax effective borrowing constraints, so marginal firms that would otherwise default are kept alive and survivors maintain more of their innovation spending. In Figure 10, the productivity gap relative to the no-shock path at year 3 is about -0.21 percent with the 15pp subsidy, -0.28 with the 5pp subsidy, and -0.33 with no intervention. The aggressive subsidy cuts the short-run loss by roughly a third. Over the first five years the ranking is monotone: larger subsidies produce smaller losses.

Figure 10. *Aggregate Productivity under Credit-Spread Subsidies: First Five Years*



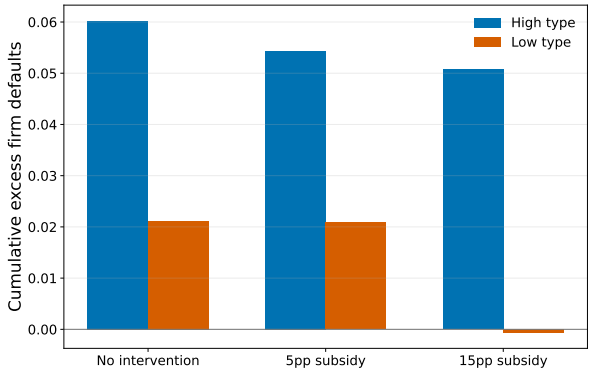
Notes: Aggregate productivity, percent deviation from the no-shock baseline path, first five years. We study interest-expense subsidies of 7.7 percent and 23.1 percent during the shock window, lump-sum financed. These are equivalent to increasing the debt tax shield by 5 and 15 percentage points, so we label them 5pp and 15pp.

What the subsidy buys: fewer defaults, almost all low-type. Figure 11 shows the composition of the prevented defaults. Panel (a) reports cumulative *excess* firm defaults (over and above the no-shock path) during the first five years, by type. With no intervention, excess defaults hit high types roughly three-to-one (0.060 versus 0.021); in totals (panel (b)), about 21 percent of high-type lines default within five years, against 13 percent of low-

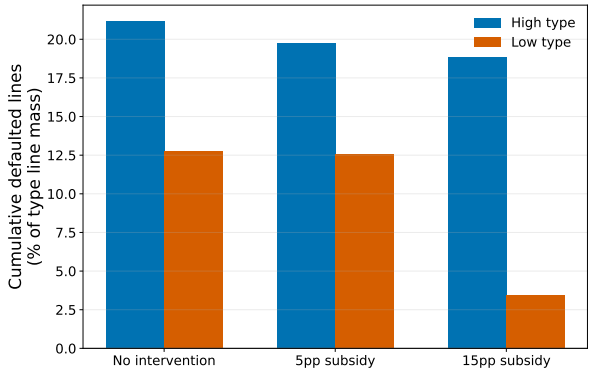
type lines. Subsidies barely dent high-type defaults: even the 15pp subsidy lowers excess high-type defaults only from 0.060 to 0.051, and total defaulted high-type lines from 21.2 to 18.9 percent of the type's line mass. It essentially *eliminates* excess low-type defaults, cutting total defaulted low-type lines from 12.7 to 3.4 percent; indeed, slightly fewer low-type firms default than with no shock at all. The marginal firm a blanket credit-spread subsidy saves is thus a low type; high-type defaults barely respond to cheaper spreads.

Figure 11. *Composition of Defaults under Credit-Spread Subsidies*

(a) *Cumulative excess firm defaults, first five years (in terms of product line mass)*



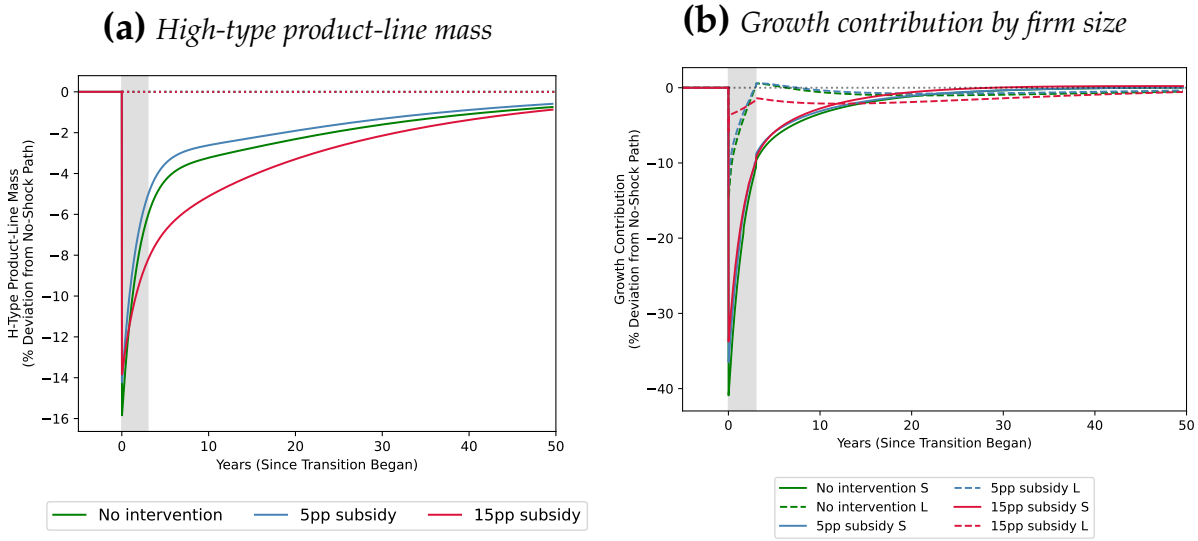
(b) *Total defaulted lines, first five years (% of own type's line mass)*



Notes: Panel (a): firm defaults in excess of the no-shock path, cumulated over five years, by type. Panel (b): total defaulted lines over the same period, including normal-times defaults, percent of each type's own line mass; the two panels use different concepts and denominators. The slightly negative 15pp low-type bar in panel (a) indicates fewer low-type defaults than with no shock at all.

Saving low types depresses reallocation to high types. Preventing low-type defaults changes who owns product lines after the shock. Figure 12(a) traces high-type firms' line mass over fifty years. The shock wipes out roughly 13–16 percent of high-type lines at impact (-15.8 percent with no intervention, -14.0 with 5pp, -13.4 with 15pp). The aggressive subsidy cushions the initial destruction but recovery is slower. With fewer low-type defaults, fewer lines migrate to high-potential owners through bankruptcy, and the 15pp economy has the lowest high-type line mass even at year 50 (about -0.9 percent below path, versus -0.7 with no intervention and -0.6 with 5pp). Panel (b) shows the counterpart in growth contributions: small firms' contribution falls roughly 34–41 percent at impact across scenarios (subsidies soften the drop), while under 15pp large firms' contribution sits mildly below path for decades.

Figure 12. *Reallocation under Credit-Spread Subsidies*

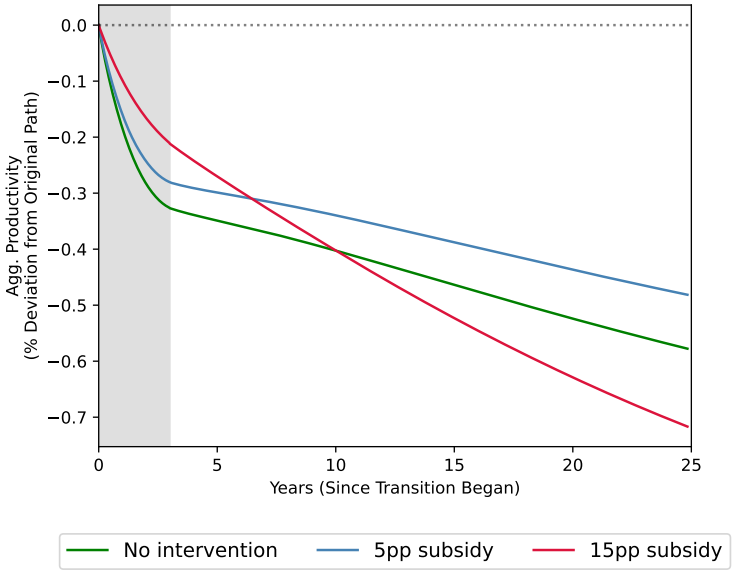


Notes: Percent deviation from the no-shock path over 50 years. Panel (a): the dotted red line marks zero; the $t = 0$ drop is the impact default wave. Panel (b): legend suffixes S/L denote small/large firms, not the low type.

The ranking reverses within a decade. The two forces, the short-run cushion and the depressed reallocation, play out at different horizons. Figure 13 extends the comparison to 25 years. The 15pp path crosses the 5pp path around year 6.5 and the no-intervention path around year 10; by year 25 it is the worst of the three, at about -0.72 percent, versus -0.58 with no intervention and -0.48 with the 5pp subsidy. The 5pp subsidy, by contrast, beats no intervention at every horizon shown. The trade-off bites only at the aggressive dose, the one large enough to stop the default wave and, with it, the reallocation on which the recovery runs. All three gaps are still widening at year 25, so these differences are lower bounds on the eventual losses rather than final values.

The experiments thus point to a tradeoff between the short-term stabilization and long-term growth. Moderate support dominates laissez-faire at every horizon shown: the 5pp subsidy stabilizes defaults and cushions the short-run loss without materially slowing reallocation toward high-potential owners. Aggressive support ends worse than no intervention within a decade: the defaults it additionally prevents are low-type, and keeping those firms alive crowds out the migration of lines to the high types that drive the recovery. Aggressive support thus recreates a zombie-lending trap, in which keeping unproductive firms alive depresses restructuring and drags on aggregate productivity (Caballero et al., 2008; McGowan et al., 2018). The subsidies we study are blanket and unanticipated: policies conditioned on firm characteristics, and the borrowing distortions

Figure 13. *The Subsidy Ranking Reverses over Longer Horizons*



Notes: Aggregate productivity, percent deviation from the no-shock baseline path, over 25 years. All gaps are still widening at year 25.

that would arise if firms anticipated intervention in normal times, are outside this experiment's scope.

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, and growth, we find that the firms that drive growth are debt-financed. High growth-potential firms innovate about four times as intensively as low-potential firms and optimally borrow more; when a credit-cost shock strikes, they suffer roughly three times the excess defaults of low-potential firms. We then analyze the transition dynamics of a credit-cost shock: an unanticipated three-year, 5-percentage-point spike in credit spreads. The shock depresses innovation, triggers a wave of bankruptcies among levered high-potential firms, and leaves aggregate

productivity persistently below its no-shock path, by about 0.46 percent at year fifteen, with no sign of the gap closing at the end of the horizon. The lasting loss is the missing contribution of small high-potential firms, whose share of aggregate growth collapses at impact and remains below its no-shock level fifteen years later. What recovery there is runs through bankruptcy reallocation and survivor selection: product lines idled in bankruptcy are re-absorbed within roughly eight years, and the firms left standing are the healthier, less-levered ones. How long that re-absorption takes matters: were bankruptcy instantaneous, the impact loss would be smaller and the recovery sharper, so the legal and search frictions in bankruptcy are a sizeable part of the productivity bill.

These dynamics discipline the reading of stabilization policy. Blanket credit-spread subsidies during the shock window cushion the short-run productivity loss, but the additional defaults an aggressive subsidy prevents are mostly low-potential firms, whose survival slows the reallocation of technologies to high-potential owners. The ranking therefore inverts over time: the aggressive 15pp subsidy delivers the smallest productivity loss inside the shock but the largest by year 25, worse than no intervention at all, while the moderate 5pp subsidy improves on no intervention at every horizon we consider. Moderate support dominates laissez-faire, but stabilizing hard enough to stop the default wave undermines the recovery it is meant to protect.

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-potential firms borrow more to finance their innovation. This means that financial shocks disproportionately hit the firms that matter most for long-run growth. Whether and how fast the economy recovers then 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; the framework is naturally suited to asking how the identity of the buyers of distressed assets (incumbents versus entrants) shapes post-crisis dynamics. 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," *The Review of Economic Studies*, 2022, *89* (1), 45–86.

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.

_ , Takeo Hoshi, and Anil K Kashyap, "Zombie Lending and Depressed Restructuring in Japan," *American Economic Review*, 2008, *98* (5), 1943–1977.

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.

Long, J. Bradford De, ""Liquidation" Cycles: Old-Fashioned Real Business Cycle Theory and the Grreat Depression," *National Bureau of Economic Research*, December 1990, pp. 1–37.

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

McGowan, Müge Adalet, Dan Andrews, and Valentine Millot, "The Walking Dead? Zombie Firms and Productivity Performance in OECD Countries," *Economic Policy*, 2018, *33* (96), 685–736.

Miao, Jianjun, "Optimal Capital Structure and Industry Dynamics," *The Journal of Finance*, December 2005, *60* (6), 2621–2659.

Osotimehin, Sophie and Francesco Pappada, "Credit Frictions and the Cleansing Effect of Recessions," *The Economic Journal*, June 2017, *127* (602), 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

$$L^S = \frac{||\Phi||}{\omega\lambda}. \tag{16}$$

A.2 Final Good Production

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

$$\begin{aligned} \ln Y &= \int_{i\in\Phi} \ln y_i di \\ &= \int_{i\in\Phi} \ln\left(\frac{q_i Y}{\lambda W}\right) di \\ &= \int_{i\in\Phi} \ln\left(\frac{q_i}{\lambda\omega}\right) di \\ &= \int_{i\in\Phi} \ln\left(q_i\right) di + \int_{i\in\Phi} \ln\left(\frac{1}{\lambda\omega}\right) di \\ &= \int_{i\in\Phi} \ln\left(q_i\right) di + \ln\left(\frac{1}{\lambda\omega}\right)\Phi \\ &= \int_{i\in\Phi} \ln\left(q_i\right) di + \ln\left(\frac{L^S}{\Phi}\right)\Phi. \end{aligned} \tag{17}$$

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

$$Y = Q\left(\frac{L^S}{\Phi}\right)^\Phi. \tag{18}$$

A.3 Normalized Firm Value

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

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

where the normalized debt drift is

$$\dot{d} = -\left[\frac{\pi n}{Y} - \chi - \frac{c(x)}{Y} - r(D,n,k)\, d\right](1-\tau_{tax}) + \frac{\varphi(p)}{Y} - g_Y d. \tag{19}$$

A.4 Firm Optimal Policies

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

$$0 = n + \left[\frac{\partial v}{\partial d}(d,n,k) + \psi_D\right][1 + 2\kappa p]\, n, \tag{20}$$

so

$$p^* = -\frac{1}{2\kappa}\left[1 + \frac{1}{\frac{\partial v}{\partial d}(d,n,k) + \psi_D}\right]. \tag{21}$$

The first-order condition for innovation intensity x is

$$0 = \left[\frac{\partial v}{\partial d}(d,n,k) + \psi_D\right](1-\tau_{tax})\eta\frac{1}{\theta_k}x^{\eta-1}n^{\eta(1-\sigma)} + n\,[v(d,n+1,k) - v(d,n,k)], \tag{22}$$

which implies

$$x^* = \left(-n^{1-\eta(1-\sigma)}\frac{\theta_k}{(1-\tau_{tax})\eta}\frac{v(d,n+1,k) - v(d,n,k)}{\frac{\partial v}{\partial d}(d,n,k) + \psi_D}\right)^{\frac{1}{\eta-1}}. \tag{23}$$

B Computational Strategy

B.1 Balanced Growth Path

Given a parameter vector, we solve for the BGP 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. (2022). 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 transition and policy experiments, we compute deterministic transition paths starting from the BGP. We guess time paths for creative destruction and for the wedge between output and consumption growth; consumption growth determines the risk-free rate through the Euler equation, while output growth enters the normalization of values and debt by Y_t. 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 approximation 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.