Firm Exit and Financial Frictions*

Gideon Bornstein[†] Laura Castillo-Martinez[‡]

June 2026

Abstract

Governments often intervene to prevent firm closures during crises, fearing that financially constrained but viable firms may fail. We develop a firm dynamics model with incomplete financial markets and show how financial frictions generate excessive firm exit. A key statistic governing this dynamic inefficiency is the marginal propensity to exit with debt. Using confidential U.S. Census data, we estimate the relationship between debt and exit and use it to discipline the model. The calibrated model implies that eliminating financial frictions reduces firm exit from 9.3% to 5.0% and generates welfare gains of 3.6% in consumption-equivalent terms. We show that the welfare costs of financial frictions rise sharply during financial crises but change little during standard productivity recessions. Finally, we compare government-guaranteed loans and grants, quantifying the trade-off between fiscal cost and effectiveness in preventing excessive exit.

*We thank Yan Bai (discussant), Luigi Bocola, Joao Gomes, Priit Jeenas (discussant), Urban Jermann, Sebnem Kalemli-Ozcan, Patrick Kehoe, Daniel Xu, and EAGLS, as well as participants at the Bank of Canada, Boston University, Carnegie Mellon University, Columbia University, Cornell University, Duke University, EIEF, the Federal Reserve Board, the Federal Reserve Banks of Chicago, Minneapolis, Philadelphia, San Francisco, and New York, Insper, the Inter-American Development Bank, PUC Rio, Rice University, Southern Methodist University, the University of Connecticut, the University of Maryland, UT Austin, Wake Forest, Wharton and various conferences for insightful comments. We also thank Catherine Gonzalez, Tanvi Jindal, and Weiting Miao, who provided excellent research assistance. Financial support from the National Science Foundation under grants SES-2116928 (Bornstein) and SES-2116375 (Castillo-Martinez) is gratefully acknowledged. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2606. (CBDRB-FY26-P2606-R12893). The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System

[†]The Wharton School, University of Pennsylvania & NBER. Email: gideonbo@wharton.upenn.edu.

[‡]Duke University, FRB Minneapolis & CEPR. Email: l.castillo-martinez@duke.edu.

1 Introduction

During COVID-19, governments worldwide deployed unprecedented resources to prevent firm closures—payroll grants, subsidized loans, tax deferrals—on the premise that financial distress, rather than weak fundamentals, was pushing viable firms toward exit. Similar interventions have accompanied major financial crises, reflecting a widespread concern that disruptions in credit markets can destroy productive firms and leave lasting scars on the economy. Yet the magnitude of this problem remains poorly understood.[1] How much firm exit is driven by financial constraints rather than underlying profitability? What are the aggregate and welfare consequences of such excess exit? And how should governments design interventions to mitigate these distortions?

To answer these questions, we combine a model of firm dynamics with new empirical evidence on the relationship between debt and exit. We first develop a model in which financial frictions prevent firms from borrowing against the full present value of future profits, generating excess firm exit. The model identifies a key informative statistic for the severity of financial frictions: the marginal propensity to exit with debt. Absent financial frictions, debt is irrelevant for exit decisions. We then use confidential Quarterly Financial Report data from the U.S. Census to estimate the relationship between debt and exit, controlling for firm characteristics. We find that debt is a strong predictor of exit conditional on productivity, size, and age: a one standard deviation increase in the debt-to-sales ratio raises the probability of exit by roughly a third, and this sensitivity doubles during recessions.

Guided by these estimates, we calibrate the model to jointly match standard firm-dynamics moments and the new debt–exit relationship. The calibrated model implies that financial frictions raise exit rates from 5.0% to 9.3% and generate welfare losses of 3.6% in consumption-equivalent terms. Finally, we use the calibrated model to study recessions and policy interventions. Excess exit barely moves during a standard productivity-driven downturn but rises sharply during a financial crisis, as collapsing borrowing capacity pushes out viable but indebted firms. Among the policy responses we consider, government-guaranteed loans curb excess exit at a lower fiscal cost than grants, despite generating debt overhang.

Our model embeds heterogeneous firms, endogenous entry and exit, and incomplete financial markets in general equilibrium. A representative household consumes a CES aggregate of differentiated varieties, each produced by a single firm under monopolistic competition. Firms differ in their idiosyncratic productivity, operate a linear technology in labor, and face stochastic fixed operating costs each period. When current profits fall short of these costs, firms can borrow from competitive financial intermediaries using one-period non-contingent debt. A firm defaults when its repayment obligations exceed the sum of current profits and available borrowing capacity. In the event of default, financial intermediaries take over the firm and sell it to new owners, who operate it at a fraction $(1-\kappa)$ of its original productivity. Default and exit are therefore distinct: a defaulting firm continues under

[1] In part, this is because workhorse models are silent on this margin. In the canonical firm dynamics framework of Hopenhayn (1992) and its descendants, exit is governed by productivity alone, and financial health plays no role in survival. Models that add financial frictions typically retain this property: in Cooley and Quadrini (2001) and Khan and Thomas (2013), exit is exogenous; in Arellano, Bai, and Kehoe (2019) and Ottonello and Winberry (2020), default-driven exit is endogenous, but the mass of entrants is set equal to the mass of exiting firms, so the number of operating firms never moves.

new ownership when the new owner finds operation profitable, and exits otherwise. The parameter $\kappa \in [0, 1]$ governs the severity of financial frictions: a larger productivity loss reduces lenders' recovery values, tightens borrowing limits, and, as we show next, widens the scope for excess exit.

The inefficiency generated by these frictions is purely dynamic. To establish this, we compare the decentralized allocation with the one a social planner would choose. The comparison rules out two distortions that could otherwise arise. First, there is no misallocation across operating firms: CES preferences imply uniform markups, so labor is allocated efficiently among active producers. Second, there is no static wedge at the exit margin: with CES preferences and fixed operating costs denominated in labor, a firm's flow profit equals its flow contribution to society, echoing Dhingra and Morrow (2019). The only potential distortion is dynamic, a wedge between a firm's borrowing capacity and the present value of its future profits. When $\kappa = 0$, even this wedge vanishes. Debt that defaults with certainty is economically equivalent to equity: lenders take over the firm and capture its entire future profit stream, so borrowing capacity equals the firm's full continuation value and the decentralized exit threshold coincides with the efficient one, despite incomplete markets. A highly leveraged firm may default, but it does not exit. Exit is determined solely by fundamentals.

When $\kappa > 0$, the dynamic wedge opens up. Default reduces the firm's productivity under new ownership, which lowers lenders' recovery values and tightens borrowing limits. A firm hit by an adverse shock may then be unable to borrow enough to cover its obligations and exit the economy, even though a social planner would keep it operating. Since κ is not directly observable, quantifying this distortion requires an observable counterpart, and the model delivers one: the marginal propensity to exit with debt. Absent financial frictions, this statistic is zero: a defaulting firm continues under new ownership at full productivity, so its survival does not depend on its balance sheet. As frictions deepen, defaulted firms become less likely to continue operating, so default increasingly means exit. Because higher debt makes default more likely, debt becomes a stronger predictor of exit. We prove that, under mild regularity conditions, the marginal propensity to exit with debt is increasing in κ in a two-period version of the model, and confirm that the same monotonic relationship holds in the full quantitative model. This result motivates our empirical strategy: the sensitivity of exit to indebtedness, conditional on firm fundamentals, provides a direct measure of the severity of financial frictions.

Estimating this sensitivity requires linking firms' balance sheets to their survival. We combine the Quarterly Financial Reports (QFR), which provide detailed balance sheet and income statement information for a representative sample of manufacturing firms, with the Longitudinal Business Database (LBD), which tracks entry and exit for the universe of U.S. employers. Unlike datasets restricted to publicly listed companies, this combination covers small and medium-sized firms, where most exit occurs. Our measure of indebtedness is the ratio of net short-term debt to sales, mirroring the short-term borrowing that covers operating shortfalls in the model. We find that this ratio is a strong predictor of exit conditional on productivity, size, age, and other fundamentals: a one standard deviation increase raises the probability of exit by 1 percentage point, a 32% increase relative to the baseline exit rate. The relationship is stable across firm sizes and age groups but strengthens significantly during recessions, with the marginal effect of indebtedness roughly doubling. We complement the U.S. analysis with Italian firm-level data from Orbis and find qualitatively similar but quantitatively larger effects, consistent with deeper financial frictions in Italy.

We then turn to the quantitative analysis. We calibrate the model to U.S. data, targeting standard moments of the firm size and age distribution used in the firm dynamics literature: age-specific exit rates, startup size, and growth profiles by initial size. The only additional moment we introduce is the debt–exit regression coefficient, which provides the empirical discipline for κ. The calibrated value, $\kappa = 0.25$, implies that new owners operate defaulted firms at 75% of their original productivity. The calibration also matches an untargeted moment that speaks directly to κ: the model-implied recovery rate on defaulted loans is 43%, close to the 44% average recovery rate on unsecured bank loans to U.S. businesses.

The aggregate costs of financial frictions are large. We quantify them by comparing the calibrated economy to the frictionless benchmark in which $\kappa = 0$. Eliminating financial frictions reduces the exit rate from 9.3% to 5.0% and raises the mass of operating firms by 55%. Average firm size falls, but the increase in available varieties more than compensates, raising steady-state consumption by 8.3%. The welfare gain, computed along the full transition path, is 3.6% in consumption-equivalent terms. We decompose these gains into two channels: the exit channel, through which financial frictions reduce the number of operating varieties, and the productivity loss channel, through which firms that default but survive operate at reduced efficiency. The exit channel accounts for roughly three quarters of the steady-state consumption gap. The modest contribution of the productivity loss channel reflects that few defaulting firms continue operating: only 1.6% of firms default and survive under new ownership each period.

Excess exit is concentrated in a specific region of the state space. High-productivity firms rarely exit regardless of their financial position, because their borrowing capacity is large enough to absorb adverse operating cost shocks. Very low-productivity firms exit under both regimes, because their fundamentals do not justify continued operation. The firms that financial frictions push out are those in between: productive enough to survive absent debt overhang, but too indebted to cover their obligations when hit by a large cost shock. The welfare costs are therefore driven not by the most distressed firms, but by viable firms in the middle whose survival depends on access to credit. This region of the state space is also where financial crises and policy interventions operate.

We then study the economy's response to two recessionary shocks: a standard downturn, modeled as a temporary decline in aggregate TFP, and a financial crisis, modeled as a temporary rise in κ, which tightens borrowing capacity and raises the sensitivity of exit to debt directly. The two experiments deliver a sharp contrast. Under the productivity shock, exit rises by only 27 basis points on impact. The shock does not operate through the exit margin, so financial frictions play no amplification role. This is consistent with the well-documented acyclicality of exit rates across most U.S. postwar recessions (Lee and Mukoyama, 2015; Ayres and Raveendranathan, 2023). Under the financial crisis, the deterioration in recovery values triggers a surge of exit concentrated among the indebted, moderate-productivity firms identified above. The damage is also far more persistent: even after κ reverts to its steady-state level, the stock of operating firms remains depleted because rebuilding through entry is slow, and the cumulative loss of varieties keeps output below trend for an extended period. Excess exit, largely irrelevant in ordinary downturns, becomes the dominant source of welfare losses in a financial crisis.

Finally, we use the model to evaluate two government interventions motivated by the COVID-19 policy response: a direct grant inspired by the Paycheck Protection Program (PPP) and a subsidized loan inspired by the Economic Injury Disaster Loan (EIDL) program. The grant covers a firm's cash-flow shortfall, up to a share of its prior-year wage bill. The loan extends the same amount as government-guaranteed credit at below-market rates, repaid over thirty years. Both curb excess exit, delivering welfare gains of 1.67% and 1.23% in one-year consumption-equivalent terms, respectively. The grant prevents more exit on impact, because the guaranteed-loan creates debt overhang: future repayment obligations tighten borrowing limits already in the period of implementation. Yet the loan is the more cost-effective instrument. Despite similar gross disbursements, repayments reduce its net fiscal cost to 0.29% of output, against 0.75% for the grant. Both instruments, however, face a fundamental trade-off. As the disbursement cap rises, a growing share of supported firms have negative social continuation value, so welfare gains are concave in program size and eventually turn negative.

Related Literature. Our paper relates to four strands of literature: firm dynamics with financial frictions, the empirical study of firm exit, the aggregate costs of financial frictions, and misallocation.

First, we build on the literature on firm dynamics with financial frictions and default, beginning with Cooley and Quadrini (2001).[2] Within this literature, Arellano, Bai, and Kehoe (2019) and Khan, Senga, and Thomas (2021) embed default risk in general equilibrium but emphasize other margins: the former studies how working-capital constraints shape firms' labor demand, while the latter studies how financial frictions misallocate investment across firms. We show that financial frictions distort the exit margin by forcing out viable firms, and we discipline this distortion with micro evidence on the debt–exit relationship. A related contemporaneous paper is Kochen (2025), who also finds sizable losses from inefficient exit but emphasizes the firm life cycle, with young firms exiting before learning their productivity. In our model, excess exit instead arises from the interaction of debt and adverse shocks, a channel that operates across the age distribution and is disciplined by our finding that debt predicts exit conditional on age, productivity, size, and other fundamentals.

Second, we contribute to the empirical literature on selection into firm exit. Lee and Mukoyama (2015) find that exit rates are roughly acyclical and that exiters are on average less productive than incumbents. Subsequent work studies how the composition of exiters changes during downturns, emphasizing the role of productivity-enhancing reallocation (Foster, Grim, and Haltiwanger, 2016) and financial distress (Leibovici and Wiczer, 2025). In an international context, Gourinchas, Kalemli-Özcan, Penciakova, and Sander (2025) study small- and medium-enterprise failures and policy responses during COVID-19, and Castillo-Martinez (2025) shows that selection among exiters depends on the exchange rate regime. Using linked balance-sheet and administrative data from the U.S. Census, we show that debt predicts exit even after conditioning on productivity, size, age, and other fundamentals, and that this relationship strengthens during downturns.

[2]For optimal-contracting approaches, see Albuquerque and Hopenhayn (2004) on limited enforcement and Clementi and Hopenhayn (2006) on moral hazard. We follow the incomplete-markets tradition, taking the debt contract as given.

Third, our paper speaks to the literature on the aggregate costs of financial frictions. A central question in this literature is whether financial frictions generate large TFP losses through capital misallocation. Buera, Kaboski, and Shin (2011) and Khan and Thomas (2013) find sizable losses, while Midrigan and Xu (2014) and Moll (2014) argue that self-financing limits these intensive-margin effects as firms accumulate internal funds and grow out of their borrowing constraints. Financial frictions also distort the allocation between investment and innovation (Ottonello and Winberry, 2024) and restrict international trade (Manova, 2013). We complement this body of work by quantifying a less-studied channel: financial frictions distort the number of operating firms.

Finally, we connect to the broader literature on misallocation and aggregate productivity. Restuccia and Rogerson (2008) and Hsieh and Klenow (2009) show that dispersion in firm-level wedges can generate large TFP losses. Dhingra and Morrow (2019) show that under CES preferences, monopolistic competition with heterogeneous firms delivers the efficient number of firms and the efficient allocation across them. These frameworks are static. Our model inherits their static efficiency properties but identifies a dynamic source of inefficiency through the extensive margin: financial frictions drive a wedge between firms' social continuation value and their borrowing capacity, generating excess exit.

Outline. Section 2 develops the model. Section 3 establishes the efficiency properties of the model and links the sensitivity of exit to debt to the degree of financial frictions. Section 4 uses firm-level data from the U.S. Census Bureau to estimate the sensitivity of exit to indebtedness, conditional on firm fundamentals. Section 5 calibrates the model to U.S. data, quantifies the aggregate costs of financial frictions, studies the economy's response to aggregate shocks, and evaluates government interventions. Section 6 concludes.

2 Model

We start with a general equilibrium model featuring firm dynamics and imperfect financial markets. This section details the environment, characterizes the agents' optimization problems, and provides a definition of the equilibrium.

2.1 Consumers

The economy is populated by a representative household that supplies its unit of labor inelastically and receives the real wage, w_t. It has preferences over consumption represented by the expected utility function

$$U = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \ln \left(C_t\right),$$

where β is the discount factor. The aggregate consumption good is a Dixit-Stiglitz aggregator of differentiated varieties, $C_t = \left[\int_0^{J_t} c_{jt}^{\frac{\epsilon-1}{\epsilon}} dj\right]^{\frac{\epsilon}{\epsilon-1}}$, where c_{jt} represents consumption of variety j at time t, ϵ governs the elasticity of substitution across varieties, and J_t is the measure of available varieties at time t. Since each variety is produced by a single firm, J_t is also the measure of operating firms.

The household owns all firms in the economy. Thus, in addition to labor income, it receives net profits generated by production at time t. In order to smooth consumption, the household trades in risk-free bonds. We normalize the aggregate price index to one, so that inter-temporal optimization delivers a standard Euler equation:

$$\frac{1}{1+r_t} = \beta \mathbb{E}_t \left[\frac{C_t}{C_{t+1}} \right] \equiv \mathbb{E}_t \mathcal{M}_{t,t+1} \,,$$

where $\mathcal{M}_{t,t+1}$ is the stochastic discount factor.

2.2 Firms

Firms are monopolistically competitive. They operate a linear technology in labor and differ in their productivity levels, $y_{jt} = Z_t a_{jt} n_{jt}$. Z_t denotes common productivity, a_{jt} denotes idiosyncratic productivity of firm j, and n_{jt} denotes the amount of labor firm j is using in production, all at time t. We assume that idiosyncratic productivity follows a Markov process with transition function given by $H(a_{jt+1}|a_{jt})$.

Each period, an arbitrarily large mass $\mathcal{M}$ of potential entrants draw an idiosyncratic entry cost from a uniform distribution on $[0, \mathcal{M} f_e]$. Potential entrants observe their cost draw before deciding whether to enter, but must pay it before observing their productivity, so entry is an irreversible commitment. Entry therefore follows a threshold rule: a potential entrant pays if and only if its draw falls below an endogenous cutoff. Entry costs are expressed in effective units of labor, so an entrant drawing cost x pays $\frac{x w_t}{Z_t}$ at time t. Under this formulation, total entry costs in effective unit of labor equal $\frac{1}{2} f_e M_e^2$, where M_e is the endogenous measure of entrants.[3] Upon entry, each firm draws its initial productivity from distribution $F(a)$.

Both incumbents and entrants are required to pay a fixed operating cost, f_{jt}, also in effective units of labor. The level of fixed costs, f_{jt}, is stochastic and follows a distribution with cumulative distribution function $G(f)$. A firm can choose not to pay this cost and exit the market. Thus, in order to continue operation, the firm j must pay $\frac{f_{jt} w_t}{Z_t a_{jt}}$.

Firms have access to one-period debt securities that can help pay off the fixed operating costs. We denote by b_{jt} the consumption goods firm j promises to repay in period $t+1$, in exchange for $q_{jt} b_{jt}$ today, where q_{jt} is the endogenous price of these securities. We allow firm owners to renege on their previous debt obligations after observing their productivity levels. In the event of default, financial intermediaries take over the firm and the firm owner walks away with nothing. We restrict $b_{jt} \geq 0$, so that firms are not allowed to accumulate cash holdings.

[3] To see this is the case, let the threshold level below which entrants pay the entry cost be $\bar{f}_e$. Then, total entry costs are

$$\mathcal{M} \int_0^{\bar{f}_e} \frac{x}{\mathcal{M} f_e} dx = \frac{1}{f_e} \left(\frac{\bar{f}_e^{\,2}}{2} \right) = \frac{1}{2} f_e M_e^2,$$

where $M_e \equiv \bar{f}_e / f_e$ is the endogenous measure of entrants that pay for entry. Note that $\mathcal{M}$ has no effect on equilibrium outcomes as long as it is large enough.

Conditional on operating, the firm sets prices and produces. Flow profits are given by

$$\pi_{jt} = \max_{\{p_{jt}, y_{jt}\}} \quad p_{jt} y_{jt} - \frac{y_{jt} w_t}{Z_t a_{jt}},$$
$$\text{s.t.} \quad y_{jt} = (p_{jt})^{-\epsilon} Y_t,$$

where p_{jt} and y_{jt} denote the price and quantity of firm j, respectively, and Y_t denotes aggregate production.

Solving the firm's problem, we obtain that profits are given by

$$\pi_t(a_{jt}) = \frac{1}{\epsilon} a_{jt}^{\epsilon-1} \overline{A}_t^{2-\epsilon} N_t Z_t, \tag{2.1}$$

where N_t is the measure of aggregate labor employed in direct production and $\overline{A}_t = \left(\int_0^{J_t} a_j^{\epsilon-1} dj\right)^{\frac{1}{\epsilon-1}}$ is a weighted average of operating firms' idiosyncratic productivity.[4]

At the end of the period, firms rebate all excess profits to their owners as dividends, d_{jt}, which we restrict to be positive, ruling out equity financing:

$$d_{jt} = \pi_{jt} - \frac{f_{jt} w_t}{Z_t a_{jt}} - b_{jt-1} + q_{jt} b_{jt} \geq 0. \tag{2.2}$$

It is convenient for the recursive formulation to define repayment obligations at the beginning of the period as the sum of debt holdings and fixed operating costs, $x_{jt} = b_{jt-1} + \frac{f_{jt} w_t}{Z_t a_{jt}}$. The idiosyncratic state of a firm is its current idiosyncratic productivity a_{jt} and its repayment obligations x_{jt}, whereas the aggregate state Ω_t is the current common productivity shifter Z_t and the measure Λ_t of entrants and incumbents over productivity and debt levels. Since firms with the same idiosyncratic state make the same choices, we drop the index j for the rest of this section.

Prior to formally describing the firm's recursive problem, we provide a brief outline. The value of the firm is the discounted sum of its stream of future dividends. Every period, the firm chooses whether to default, how much to borrow, and how much to pay in dividends. It is subject to a budget constraint and a non-negative dividend condition. If the firm can borrow to cover its financial obligations, it always chooses to do so. This option dominates defaulting as the non-negative dividend condition ensures the firm's value is weakly positive. It then follows that the firm defaults only when its budget set is empty.

The Firm's Recursive Problem. To derive the default policy function, let $\bar{L}(a_t; \Omega_t)$ denote the endogenous borrowing limit of a firm with productivity a_t when the aggregate state of the economy is Ω_t. The default decision boils down to whether the sum of profits and this borrowing limit suffices to cover the firm's repayment obligations. Let $\bar{f}_t = \bar{f}(a_t, b_{t-1}; \Omega_t)$ be the highest fixed operating cost that ensures the non-negative dividend condition is satisfied. From condition (2.2), this threshold level is

$$\frac{\bar{f}_t w_t}{Z_t a_t} = \pi_t - b_{t-1} + \bar{L}(a_t; \Omega_t). \tag{2.3}$$

[4] See Appendix A.1 for full derivation.

The firm chooses to repay if its fixed operating cost is below the threshold, $f_t \leq \bar{f}_t$, and defaults otherwise.

We are now ready to consider the problem of an incumbent firm. Let $V(a_t, x_t; \Omega_t)$ denote the discounted value of the firm after all period t shocks are realized. The value of the firm is equal to zero for any state $(a_t, x_t; \Omega_t)$ such that the budget set is empty. For all other states, firms choose new borrowing b_t and dividends d_t to solve

$$V(a_t, x_t; \Omega_t) = \max_{\{b_t \geq 0, d_t \geq 0\}} d_t + \mathbb{E}_t \left[\mathcal{M}_{t,t+1} \widetilde{V} \left(a_{t+1}, b_t + \frac{f_{t+1} w_{t+1}}{Z_{t+1} a_{t+1}}; \Omega_{t+1} \right) \right], \tag{2.4}$$
$$\text{s.t.} \quad d_t = \pi(a_t; \Omega_t) - x_t + q(a_t, b_t; \Omega_t) b_t,$$

where $\pi(a_t; \Omega_t)$ is given by equation (2.1), $q(a_t, b_t; \Omega_t)$ denotes the debt pricing schedule, which we derive in the next subsection, and $\mathcal{M}_{t,t+1}$ is the stochastic discount factor. The future value of the firm is given by

$$\widetilde{V}(a_{t+1}, x_{t+1}; \Omega_{t+1}) = \begin{cases} V(a_{t+1}, x_{t+1}; \Omega_{t+1}) & \text{if } f_{t+1} \leq \bar{f}_{t+1}, \\ 0 & \text{otherwise.} \end{cases}$$

2.3 Financial Sector

Firms borrow from competitive financial intermediaries that diversify their lending portfolios and thus avoid exposure to idiosyncratic risk. This implies that, absent aggregate uncertainty, the expected return on a loan to a firm should be equal to the real return on risk-free debt, $1 + r_t$.

In case of default, financial intermediaries sell the firm to a new owner. Under new ownership, firms operate with lower productivity $(1 - \kappa)a$, where a is the original productivity level. This assumption captures the idea that the firm owner has non-transferable knowledge on how to run the firm. Similar to models of financial frictions where a fraction of the capital stock vanishes upon default, here a fraction of the productive knowledge is lost. As we will discuss in the next section, the parameter $\kappa \in [0, 1]$ governs the degree of financial frictions in the model.

In order for financial intermediaries to break even on lending, the debt pricing schedule solves the following equation for $b_t > 0$,

$$q(a_t, b_t; \Omega_t) = \mathbb{E}_t \left[\mathcal{M}_{t,t+1} G\left(\bar{f}(a_{t+1}, b_t; \Omega_{t+1}) \right) \right]$$
$$+ \mathbb{E}_t \left[\mathcal{M}_{t,t+1} \frac{\widetilde{V}\left((1-\kappa) a_{t+1}, \frac{f_{t+1} w_{t+1}}{Z_{t+1}(1-\kappa) a_{t+1}}; \Omega_{t+1} \right)}{b_t} \mathbb{1}\left[f_{t+1} > \bar{f}(a_{t+1}, b_t; \Omega_{t+1}) \right] \right]. \tag{2.5}$$

The first term captures the expected repayment value. Recall that the firm's default decision has a cutoff form. Repayment in period $t + 1$ occurs as long as the fixed operating cost falls below the threshold $\bar{f}_{t+1}$, that is with probability $G(\bar{f}_{t+1})$. The second term captures the expected default value. Following default, financial intermediaries take over the firm and sell it to new owners. In exchange, financial intermediaries receive the value of the firm, $\widetilde{V}\left((1-\kappa) a_{t+1}, \frac{f_{t+1} w_{t+1}}{Z_{t+1}(1-\kappa) a_{t+1}}; \Omega_{t+1} \right)$, which accounts for the productivity loss associated with new ownership and the lower repayment obligations implied by default.

Given the debt pricing schedule, the borrowing limit of a firm with productivity a_t when the aggregate state of the economy is Ω_t is given by

$$\bar{L}(a_t; \Omega_t) = \sup_{b_t > 0} q(a_t, b_t; \Omega_t) b_t.$$

2.4 Firm Exit

The timing within each period is as follows. First, potential entrants decide whether to pay the entry cost and enter the market. Second, firms observe their productivity level as well as their fixed operating costs. Based on that information, they decide whether to default on their debt obligations. Firms that default are taken over by financial intermediaries and are sold to new owners. Next, firms without any debt obligations, including the defaulted firms, decide whether to exit the economy. Third, all firms choose their prices and quantities, borrow from intermediaries, and production takes place. Finally, profits are realized, firms repay their debt obligations as well as operating costs, and dividends are distributed to owners.

Firms exit the economy if the current owner decides not to repay the debt obligations and, upon default, the new owner chooses not to pay the fixed operating costs. We define the exit threshold $f_x(a_t, b_{t-1}; \Omega_t)$ as the level of fixed operating cost above which a firm with productivity a_t and debt obligations b_{t-1} prior to potential default exits the economy. It follows that

$$f_x(a_t, b_{t-1}; \Omega_t) = \max \left\{ \bar{f}(a_t, b_{t-1}; \Omega_t), \bar{f}((1-\kappa)a_t, 0; \Omega_t) \right\}. \tag{2.6}$$

The first term in the maximum is the threshold level that triggers default. The second term is the threshold level above which new owners choose not to pay the fixed operating costs.

2.5 Stationary Equilibrium

Let us now define the stationary equilibrium without aggregate shocks, where $Z_t = 1$ for all t. For this definition, we denote the joint distribution of firms across productivity levels and debt levels by $\Lambda(a, b)$. This measure includes entrants and incumbents prior to any financial decisions and prior to the realization of fixed operating costs. That is, it does not account for exit or changes in firm ownership. Four components characterize the law of motion for this joint distribution: (i) the default and exit decision of firms, (ii) the borrowing decision of firms, (iii) the measure of entrants, and (iv) the exogenous law of motion for productivity. The law of motion for the joint distribution is defined as follows. For all Borel sets $\mathcal{A} \times \mathcal{B} \subset \mathbb{R}^{+} \times \mathbb{R}$,

$$\begin{aligned}\Lambda'(\mathcal{A}\times\mathcal{B}) = &\int_f \int \int_{a'\in\mathcal{A}} \mathbb{1}\left[b'\left(a, b+\frac{fw}{a}\right)\in\mathcal{B}\right] dH(a'|a)d\Lambda(a,b)dG(f) \\ &+\int_f \int \int_{a'\in\mathcal{A}} \mathbb{1}\left(f>\bar{f}(a,b)\right)\mathbb{1}\left[b'\left((1-\kappa)a, \frac{fw}{(1-\kappa)a}\right)\in\mathcal{B}\right] dH(a'|(1-\kappa)a)d\Lambda(a,b)dG(f)\ , \\ &+\mathbb{1}\left(0\in\mathcal{B}\right)M_e'\int_{a'\in\mathcal{A}} dF(a')\ ,\end{aligned} \tag{2.7}$$

where $b'(a, x)$ is the borrowing choice of a firm with productivity a and repayment obligations x. If a firm with state variables (a, x) chooses to default, then $b'(a, x) = \emptyset$.[5] The first term in the transition function comes from firms that deliver on their debt obligations this period. The second term captures firms that default but operate under new ownership this period. The third term corresponds to next period's entrants. Note that these new firms enter the market with no debt.

In the stationary equilibrium, the distribution $\Lambda(a, b)$ is constant over time. The definition of the stationary Markov-perfect equilibrium is as follows:

Definition **1** (Equilibrium)**.** A stationary Markov-perfect equilibrium is a set of aggregate allocations $\{\bar{A}, N\}$, aggregate real wage w, a debt pricing schedule q, policy functions, a measure of entrants M_e and a distribution of firms over productivity and debt levels, $\Lambda(a, b)$, such that:

1. Borrowing and default policies solve the firm's problem.

2. Debt pricing policy solves equation (2.5).

3. Free entry condition holds:

$$f_e M_e w = \int_f \int_a \widetilde{V}\left(a, \frac{fw}{a}\right) dF(a)dG(f).$$

4. The level of labor employed in production, N, is given by

$$\begin{aligned}N = &\int\int \mathbb{1}\left(f \le \bar{f}(a,b)\right) n(a)d\Lambda(a,b)dG(f) \\ &+\int\int \mathbb{1}\left(f>\bar{f}(a,b)\right)\mathbb{1}\left(f\le\bar{f}((1-\kappa)a,0)\right) n((1-\kappa)a)d\Lambda(a,b)dG(f),\end{aligned}$$

where $n(a)$ is the policy function indicating how many workers a firm with productivity a hires for production. Note that this policy function is independent of b. The first term is the hiring of firms that repay their debt, and the second term is the employment of firms which default but continue operating.

[5] Since $\emptyset \notin \mathcal{B}$ for any Borel set $\mathcal{B} \subset \mathbb{R}$, defaulting firms are excluded from the first term and post-default exiting firms are excluded from the second term, without requiring explicit continuation indicators.

5. Labor market clears:

$$1 = N + \int\int \mathbb{1}\left(f \le \bar{f}(a,b)\right) \frac{f}{a}\, dG(f)\, d\Lambda(a,b)$$
$$+ \int\int \mathbb{1}\left(f > \bar{f}(a,b)\right) \mathbb{1}\left(f \le \bar{f}((1-\kappa)a, 0)\right) \frac{f}{(1-\kappa)a}\, dG(f)\, d\Lambda(a,b) + \frac{f_e M_e^2}{2}.$$

6. Aggregate productivity satisfies its definition.

7. The distribution of firms is stationary.

2.6 Discussion

Firms cannot accumulate cash holdings. Allowing retained earnings would alleviate the costs of financial frictions, at the potential cost of too few indebted firms relative to the data. An alternative approach is to allow cash holdings while adding a feature that curtails firms' incentive to save — for instance, a tax advantage of debt (Jermann and Quadrini, 2012) or agency frictions that limit the cash a firm can hold (Arellano, Bai, and Kehoe, 2019). Our calibrated model matches the share of firms with a positive net debt position well (Section 5.1), and the simpler formulation delivers a cleaner theoretical benchmark: absent financial frictions, the decentralized equilibrium is efficient, a result we establish in the next section.

Upon default, financial intermediaries take over the firm and sell it to new owners, who operate it at a fraction $(1-\kappa)$ of its original productivity. We do not allow the original owner and lender to renegotiate so that the owner continues without a productivity loss. An alternative interpretation is that the owner does remain in place, but the time, attention, and legal costs of restructuring reduce effective productivity. Under either interpretation, the key implication is the same: default reduces future cash flows and tightens the endogenous borrowing limit. This connects the framework to the empirical regularity that most corporate borrowing is earnings-based: Lian and Ma (2021) show that over 80% of U.S. non-financial corporate debt is cash-flow based. In our model, borrowing limits depend on the present discounted value of future profits, providing a structural rationale for this pattern.

Entry costs are financed entirely through equity. As a result, financial frictions affect the entry margin only indirectly, through their effect on the expected present value of a new firm. Had entry required debt financing, financial frictions would directly raise the cost of entry and amplify their aggregate effects beyond the exit margin.[6] We abstract from this channel to isolate the costs of financial frictions operating through firm exit—the margin we identify theoretically in Section 3 and discipline empirically in Section 4.

The model abstracts from physical capital and capital accumulation. Introducing capital would affect the analysis in two ways that work in opposite directions. First, physical capital could serve as collateral, improving lenders' recovery values upon default. This would relax borrowing constraints and reduce excess exit. Second, with capital as a production input, financial frictions would also distort

[6] An extensive literature studies how financial frictions shape firm entry directly; see Quadrini (2009) for a survey and Cuciniello, Michelacci, and Paciello (2025) for a recent general-equilibrium quantification of this channel.

the allocation of capital across firms—an intensive-margin misallocation channel studied extensively in the literature (Buera, Kaboski, and Shin, 2011; Khan and Thomas, 2013). By abstracting from capital, we isolate the exit margin as the sole channel through which financial frictions operate in our model, making our framework complementary to this body of work.

3 The Role of Financial Frictions in Firm Exit

In this section, we examine how financial frictions shape firm exit decisions in the model. We first show that market imperfections, captured by the financial friction parameter κ, create a wedge between the efficient and decentralized exit thresholds. We then explore how the empirical relationship between debt and firm exit is informative of the degree of financial frictions.

3.1 Dynamic Misallocation

A firm exits the economy if two conditions are satisfied. First, the current owner of the firm must choose not to pay its current operating costs and default on its debt obligations. Second, the new owner must not find it profitable to operate the firm. In the previous section, we have shown that these two conditions can be stated jointly as a threshold condition on current fixed operating costs, reproduced here for convenience:

$$f_x(a_t, b_{t-1}; \Omega_t) = \max\Big\{\bar{f}(a_t, b_{t-1}; \Omega_t),\ \bar{f}((1-\kappa)a_t, 0; \Omega_t)\Big\}, \tag{3.1}$$

$$\text{where} \quad \bar{f}(a_t, b_{t-1}; \Omega_t) = a_t Z_t \left[\frac{\pi_t(a_t)}{w_t} + \frac{1}{w_t}\left(\bar{L}(a_t; \Omega_t) - b_{t-1}\right)\right]. \tag{3.2}$$

In either case, the threshold depends on two components: flow profits, which capture the current cash flow generated by production given the firm's productivity a_t, and net borrowing capacity, which summarizes the amount of external funds the firm can raise once existing debt obligations are taken into account. Together, these components determine whether the firm remains active.

To assess the allocative implications of financial frictions on firm exit, we compare the firm's private exit rule with the socially efficient one. The social planner would like a given firm to continue operating as long as the net present value of its social contribution is positive.[7] Let $W_{FB}(a_t, f_t; \Omega_t)$ denote the social value of a firm with productivity a_t and operating cost f_t that is active at time t. We show in the Appendix that the social value of a firm, in units of final output, can be written recursively as follows,

$$W_{FB}\left(a_t, f_t; \Omega_t\right) = \frac{1}{\epsilon-1}\overline{A}_t^{2-\epsilon} Z_t N_t a_t^{\epsilon-1} - \overline{A}_t \frac{f_t}{a_t} + \mathbb{E}_t \mathcal{M}_{t,t+1} \tilde{W}_{FB}\left(a_{t+1}, f_{t+1}; \Omega_{t+1}\right), \tag{3.3}$$

where $\tilde{W}_{FB}(a_{t+1}, f_{t+1}; \Omega_{t+1}) \equiv \max\{W_{FB}(a_{t+1}, f_{t+1}; \Omega_{t+1}), 0\}$ is the social value of the firm before the social planner decides whether the firm should pay next period's fixed cost. The first term in equation (3.3) represents the flow social benefit of operating the firm. Due to love of variety, the

[7] We relegate the formal definition of the efficient allocation to Appendix A.2.

benefit of any operating firm is positive and it is increasing in its idiosyncratic productivity a_t. The second term captures the flow social cost of operating the firm: the resources used in fixed operating activities that reduce available labor for direct production. The final term represents the expected social continuation value, reflecting the discounted value of future social contributions.

The efficient exit threshold $\overline{f}_x^{FB}(a_t;\Omega_t)$ is implicitly defined by the indifference condition between operating and exiting:

$$\overline{f}_x^{FB}(a_t;\Omega_t) = a_t Z_t \left[\frac{1}{\epsilon-1}\overline{A}_t^{1-\epsilon} N_t a_t^{\epsilon-1} + \frac{1}{\bar{A}_t Z_t}\mathbb{E}_t \mathcal{M}_{t,t+1}\tilde{W}_{FB}\left(a_{t+1}, f_{t+1};\Omega_{t+1}\right)\right]. \tag{3.4}$$

Comparing the decentralized and efficient thresholds illustrates how financial frictions distort exit dynamics. Two potential wedges may arise: a static wedge reflecting differences between the firm's private flow profit and its flow social contribution, and a dynamic wedge capturing discrepancies between the firm's net borrowing capacity and the net present value of its future social contributions. When both wedges are zero, the decentralized exit threshold coincides with the planner's efficient threshold. The following proposition establishes that there is no static wedge in our model.

Proposition **1.** *If $\beta = 0$, the decentralized and efficient allocations coincide.*

We relegate all proofs to Appendix A.3, but let us sketch the proof here to build intuition. When $\beta = 0$, absent debt obligations, the decentralized and efficient exit thresholds are equalized. Firms have no future value, so there is neither borrowing in the decentralized equilibrium nor a continuation social value in the planner's problem. The exit condition then depends solely on contemporaneous flow payoffs. Solving for the equilibrium wage and substituting into the definition of flow profits in equation (2.1) shows that the firm's private flow profit coincides with its social flow contribution.[8]

This result hinges on three key modeling assumptions: CES preferences, exogenous aggregate labor supply, and fixed operating costs expressed in labor units. Intuitively, there is no misallocation of production inputs because markups are identical across firms, so labor is efficiently allocated across the set of active producers. Although an aggregate markup could, in principle, distort total labor supply as in Edmond, Midrigan, and Xu (2023), this channel is also absent because labor supply is fixed. Moreover, since fixed operating costs are denominated in labor units, they do not inherit the markup distortions embedded in output prices. This result echoes the main insight in Dhingra and Morrow (2019), namely that in static environments with these same features, productivity heterogeneity affects profits but not the efficiency of resource allocation.

With static misallocation ruled out, we now turn to the inter-temporal dimension of the model. In general, the decentralized and efficient exit thresholds differ because a firm's borrowing capacity is not equal to the discounted value of its future social contribution. The wedge arises from the difference in the second term of equations (3.2) and (3.4), which captures the dynamic component of the decision. Financial frictions limit the firm's ability to borrow against future profits, preventing it from fully capitalizing its potential future value.

[8] Both thresholds are defined conditional on the aggregate state Ω_t. The remainder of the proof shows that the aggregate state Ω_t, which includes the distribution of incumbents and entrants before the realization of fixed operating costs, is identical across the decentralized and efficient allocations.

Absent financial frictions ($\kappa = 0$), the model continues to feature incomplete markets in the sense that debt obligations are non-state-dependent. Nonetheless, allowing financial intermediaries to take over the firm in the event of default provides sufficient financial spanning to replicate the planner's allocation. The following proposition formalizes this result.

PROPOSITION **2.** *If $\kappa = 0$, the decentralized and efficient allocations coincide.*

The proof of Proposition 2 shows that when $\kappa = 0$, the firm's borrowing limit equals the expected discounted stream of its future net profits. Since default entails no productivity loss, outstanding debt obligations do not affect a firm's exit probability. Although a highly leveraged firm may default, it can continue operating under new ownership. As a result, the threshold for exit in the decentralized equilibrium coincides with the efficient one.

Even in the presence of incomplete markets, firms can raise funds equal to their full discounted stream of net profits by issuing debt that is certain to default. In this limit, debt behaves as equity: financial intermediaries anticipate default and value their claims accordingly. Since default does not destroy productivity when $\kappa = 0$, creditors' payoffs are fully state-contingent ex post. Hence, despite the absence of explicit state-contingent contracts, the equilibrium replicates the complete-markets allocation.

To conclude, we have shown in this subsection that financial frictions in our model lead to distortions in the firm exit margin. The source of misallocation is dynamic, driven by the firm's borrowing capacity being lower than the net present value of its future profits. We next turn from normative to positive analysis, studying how financial frictions shape the relationship between firm exit and debt.

3.2 Financial Frictions and the Relationship between Debt and Exit

Our ultimate goal is to quantify the degree of dynamic misallocation in the economy. While the degree of financial frictions is not observable, we show that the marginal propensity to exit with debt provides an informative empirical moment for estimating this parameter. To build intuition and derive analytical results, we focus on a simplified two-period version of the full model.

Suppose the economy is populated by a unit measure of potential firms that differ in their initial debt levels, b_0, and in the realization of their fixed operating cost, f, drawn from distribution $G(f)$. Each operating firm generates a constant gross profit π in every period in which it operates without default. Firms borrow from competitive lenders only in the first period, b_1, according to the debt pricing schedule derived in equation (2.5):

$$q(b_1) = \beta G\left(\pi - b_1\right) + \beta \int_{\pi - b_1}^{(1-\kappa)\pi} \left[\frac{(1-\kappa)\pi - f}{b_1}\right]^{+} dG(f) \ , \tag{3.5}$$

which accounts for the assumption that, upon default in the second period, the bank can operate the firm directly but only generate $(1-\kappa)\pi$. If a firm defaults in the first period, any new debt is then priced according to:

$$q^{\kappa}(b_1) = \beta G\left((1-\kappa)\pi - b_1\right) + \beta \int_{(1-\kappa)\pi - b_1}^{(1-\kappa)^2\pi} \left[\frac{(1-\kappa)^2\pi - f}{b_1}\right]^{+} dG(f) \, . \tag{3.6}$$

Intuitively, the price of debt falls after default because lenders anticipate a lower expected return: repayment in the second period becomes less likely, as the firm's profitability has deteriorated, and the recovery value in case of another default is smaller. The corresponding borrowing limits are then $\bar{L} = \max_{b_1} q(b_1)b_1$ and $\bar{L}^\kappa = \max_{b_1} q^\kappa(b_1)b_1$, respectively.

Given these borrowing limits, the probability that a firm with initial debt b_0 exits in the first period is

$$P_x(b_0) = 1 - G\left(\max\left\{\pi + \bar{L} - b_0\,, (1-\kappa)\pi + \bar{L}^\kappa\right\}\right). \tag{3.7}$$

This expression links the probability of exit to the exit threshold defined in equation (2.6): a firm exits when its realized fixed operating cost exceeds the relevant cutoff—either the default or the exit-upon-default threshold, whichever is higher.

The following lemma characterizes the marginal propensity to exit with debt, $P_x'(b_0)$.

Lemma **1.** *The marginal propensity to exit with debt is*

$$P_x'(b) = \begin{cases} g(\pi + \bar{L} - b) & \textit{for } b < \bar{b} \\ 0 & \textit{for } b > \bar{b} \end{cases} \tag{3.8}$$

where

$$\bar{b} = \kappa\pi + \bar{L} - \bar{L}^\kappa.$$

The cutoff debt level $\bar{b}$ is the point below which the exit margin is determined by the default threshold, and above which it is determined by the exit-upon-default threshold. When debt is sufficiently low, default is always accompanied by exit. A higher level of debt increases default probability and, as a result, also the exit probability. At the cutoff debt level $\bar{b}$ itself, P_x is kinked: the left derivative is strictly positive while the right derivative is zero. For firms with debt above $\bar{b}$, exit is determined by the exit-upon-default threshold, which is independent of the level of debt. In this region, a higher level of debt raises the default probability but has no effect on the exit probability.

An immediate implication of Lemma 1 is that absent financial frictions ($\kappa = 0$), the marginal propensity to exit with debt is equal to zero. Debt is not correlated with firm exit because if running the firm is profitable, it would be sold to new owners upon default, and the firm would not exit the economy. Finally, we examine how the relationship between debt and exit responds to changes in the degree of financial frictions.

Proposition **3.** *With a decreasing density of fixed costs $g(f)$, a rise in κ (more financial frictions) weakly increases the marginal propensity to exit for any debt level, b.*

The proof of Proposition 3 shows that a higher κ expands the range of debt levels with a positive marginal propensity to exit, and that within this range $P_x'(b)$ increases with κ.[9] The intuition is straightforward. As financial frictions rise, it is less likely that defaulted firms can continue operating, since lenders recover less and refinancing terms worsen. Consequently, the default threshold becomes the dominant determinant of exit. Debt, which affects exit only through its impact on default

[9] A decreasing $g(f)$ is sufficient but not necessary. Under a log-normal distribution the proposition continues to hold as long as the exit threshold lies to the right of the mode, where the density is decreasing.

probabilities, therefore exerts a stronger influence on firms' survival decisions, raising the marginal propensity to exit with debt.

Lemma 1 and Proposition 3 show that the degree of financial frictions in the economy is tightly linked to the marginal propensity to exit with debt. In Section 5, we find that the same pattern holds also in our full calibrated model—more financial frictions increase the marginal propensity to exit with debt. Guided by this relationship, a key moment that will help us pin down the degree of financial frictions in the model is the marginal propensity to exit with debt in the data.

4 Empirical Evidence

In this section, we use firm-level data to study the relationship between debt and exit. We begin by describing the data sources for our baseline analysis, which relies on confidential microdata from the U.S. Census Bureau. We then estimate the marginal propensity to exit with respect to debt, exploring heterogeneity by firm characteristics and over the business cycle. Finally, we complement the analysis using firm-level data from Italy.

4.1 Data

We exploit two main confidential datasets provided by the U.S. Census Bureau to construct a representative sample of firm dynamics with a financial overview: the Quarterly Financial Reports (QFR) and the Longitudinal Business Database (LBD). We complement these data with the Census of Manufacturers (CM), which provides additional information on firms' production characteristics.

The QFR, perhaps the least familiar of the two, is a long-standing survey of firms in mining, manufacturing, wholesale and retail trade, and, more recently, services. Although the survey predates its administration by the Census, it has been managed there since 1982 and provides quarterly income and balance sheet statements for a stratified random sample of firms. Stratification is based on firm size: large firms with more than $250 million in book assets are sampled with certainty, while smaller firms, those with assets between $250 thousand and $250 million, are drawn from a rotating panel and observed for eight consecutive quarters. We focus exclusively on the manufacturing sector, where the coverage of small firms is broadest.[10] As noted by Crouzet and Mehrotra (2020), manufacturing coverage ranges between 5 and 8 percent depending on the quarter. To ensure representativeness, the Census Bureau assigns sampling weights to each observation, which we use throughout the analysis.

Linking the QFR micro files across quarters is not straightforward. To construct a usable panel dataset, we closely follow the guidance of Crouzet and Mehrotra (2020), who were the first to undertake this task.[11] Specifically, we link firms primarily through their employer identification number (EIN), and when the EIN is unavailable, we match firms by name and headquarters location. For the years 1994–2000, when Census sampling weights are missing, we adopt their imputed weights. We also drop miscoded observations and restrict the sample to firms with strictly positive assets and internally

[10] Outside manufacturing, only firms with assets between $50 million and $250 million are included in the rotating panel.

[11] We are grateful to Nicolas Crouzet for generously helping us work through the details of the procedure.

consistent balance sheet data. For further details on the validation of the cleaned dataset, we refer the interested reader to Crouzet and Mehrotra (2020). In our implementation, we extend the panel through 2017. The resulting sample contains about 814,000 firm–quarter observations between 1982:III and 2017:IV. While the dataset is constructed at the quarterly frequency, we consolidate it to the annual level, which is the time unit of analysis in what follows.[12]

The QFR does not contain information on firm exit. To address this limitation, we merge the QFR with the Longitudinal Business Database (LBD), which covers the universe of U.S. establishments with employees. The LBD records the year in which each establishment first appears, as well as the year it is last observed if it subsequently exits. We define firm entry as the year in which a firm's oldest establishment is born, and firm exit as the year in which its last establishment closes. Even though legal restrictions allow QFR data to be merged with other Census business datasets only in Economic Census years (historically, years ending in 2 and 7), entry and exit dates can be recovered for 90% of the firms in our QFR sample.

Beyond exit, the LBD provides information on employment and firm demographics, while output is obtained from the Census of Manufacturers (CM). These datasets allow us to complement the QFR's financial statements with additional non-financial firm characteristics. As discussed above, these variables are observed only every five years and are used as controls whenever available.

Descriptive Statistics. Table 1 reports summary statistics on real and financial characteristics of the firms in our main sample. We also present statistics for the subset of years ending in 2 or 7, the Census years, to confirm overall characteristics remain stable across this subsample. Firm fundamentals include age, sales, sales growth, capital, employment, and labor productivity. Capital is measured as net property, plant, and equipment (PP&E), corresponding to the book value of tangible fixed assets net of accumulated depreciation. Sales and capital are expressed in millions of constant 2017 U.S. dollars. Employment corresponds to the number of employees, sales growth is defined as the year-to-year percent change, and labor productivity is measured as the ratio of real output to employment.

In addition to firm fundamentals, we also report financial characteristics. While the analysis focuses on short-term financing, we present total assets and liabilities alongside a breakdown into current and non-current components for completeness. On the asset side, this includes cash holdings, trade receivables, and non-current assets. On the liability side, we distinguish between short-term bank debt, trade payables, and non-current liabilities. All financial variables are expressed in millions of constant 2017 U.S. dollars.

4.2 Debt and the propensity to exit

We now examine the relationship between debt and exit. Our goal is to estimate the conditional correlation between debt and firm exit, controlling for non-financial fundamentals. As shown in Section 3,

[12] For balance sheet stocks, we take the average of the four quarterly observations within the year; for income statement flows, we sum across quarters.

Table 1: Summary Statistics

	Full sample					Census sample				
	Mean	St. dev.	p25	p75	N	Mean	St. dev.	p25	p75	N
Firm fundamentals										
Age	16.86	9.35	10.00	26.00	119,000	17.84	9.75	11.00	29.00	19,000
Sales	17.22	126.80	4.45	91.19	119,000	10.85	81.26	3.08	42.86	19,000
Sales growth	12.52	83.39	−38.85	55.64	119,000	4.57	83.88	−52.53	48.83	19,000
Capital stock	3.71	29.29	0.65	20.03	119,000	2.40	19.34	0.43	10.17	19,000
Employment						52.73	322.10	24.00	200.00	19,000
Labor productivity						0.14	0.59	0.07	0.19	15,500
Financial characteristics										
Assets	13.88	116.10	3.53	78.36	119,000	8.86	76.92	2.56	37.45	19,000
Cash	0.83	5.13	0.09	2.75	119,000	0.69	4.16	0.08	2.01	19,000
Accounts receivable	2.52	14.58	0.81	15.64	119,000	1.72	9.63	0.58	8.17	19,000
Non-current assets	2.59	40.44	0.01	4.16	119,000	1.35	26.58	0.00	1.66	19,000
Liabilities	7.28	65.43	1.37	35.36	119,000	4.48	42.09	0.94	17.09	19,000
Short-term bank debt	0.88	5.20	0.03	3.45	119,000	0.64	3.64	0.01	2.05	19,000
Accounts payable	1.38	9.54	0.33	7.66	119,000	0.94	6.53	0.21	4.12	19,000
Non-current liabilities	3.57	39.21	0.15	11.27	119,000	2.05	25.17	0.07	5.22	19,000
Exit rate	3.17	17.52	0.00	0.00	119,000	3.70	18.88	0.00	0.00	19,000

Notes: This table reports summary statistics on firm fundamentals and financial characteristics. Sales, capital, and all financial variables are expressed in millions of constant 2017 U.S. dollars and are winsorized at the 5–95% level each year. Sales growth is the year-over-year percent change in sales. Employment corresponds to the number of employees, and labor productivity is measured as the ratio of output in millions of constant 2017 dollars to employment. The Census sample refers to years ending in 2 or 7.

this correlation is tightly linked to the degree of financial frictions in the model and will serve as a key moment for calibration.

Our main empirical measure of financial health is the ratio of net short-term debt to sales, chosen for consistency with our theoretical framework. Assets are absent from the model, while sales represent the relevant cashflows against which debt obligations must be serviced. Our emphasis on short-term debt also reflects the fact that the model explicitly features short-term borrowing. Finally, we use the net debt position to capture financial needs that cannot be met with retained earnings. Net short-term debt is defined as short-term loans plus accounts payable minus cash holdings and accounts receivable.[13]

To isolate the role of debt, we control for non-financial firm characteristics that are known to predict exit. Table A.1 in the appendix compares exiting firms to incumbents in the year preceding exit and confirms well-known patterns: exiters are younger, smaller, have lower sales growth, hold less capital, and feature lower labor productivity. These variables enter the regressions as controls and, as reported in the appendix, operate in the expected directions.

We proceed in three steps. We begin with a baseline specification that relates firm exit to debt while controlling for fundamentals. We then explore how this relationship varies with firm characteristics by interacting debt with size and age. Finally, we investigate whether the link between debt and exit changes systematically over the business cycle.

[13] All main results are robust to using the debt-to-assets ratio as an alternative measure of leverage; see Table A.2 for details.

Baseline regression. To start, our empirical model is given by

$$Exit_{ist} \;=\; \alpha\, DebtRatio_{is,t-1} \;+\; X'_{is,t-1}\,\omega \;+\; \delta_{st} \;+\; \varepsilon_{ist}, \tag{4.1}$$

where i indexes firms, s denotes three-digit industries, and t is the year. The dependent variable $Exit_{ist}$ is a firm-level indicator equal to one if the firm exits the market in year t. The key explanatory variable, $DebtRatio_{is,t-1}$, is the net short-term debt-to-sales ratio measured in year $t-1$. The vector $X_{is,t-1}$ collects non-financial firm characteristics measured in $t-1$, and δ_{st} are industry–year fixed effects. The timing mirrors the theoretical framework such that future propensity to exit correlates with current firm characteristics. Standard errors are clustered by industry-year, since some explanatory variables vary only at that level.

Non-financial controls are grouped into *main* and *additional* firm fundamentals. The *main* set comprises (i) age, captured by a dummy $Young_{is,t-1}$ equal to one if the firm is five years old or younger; (ii) size, captured by firm-size fixed effects based on sales quintiles constructed at the industry–year level; (iii) sales growth; and (iv) the log of the capital stock. The *additional* set comprises the log of employment and the log of labor productivity. This distinction between main and additional controls is driven by data availability, as discussed above.

The main results are reported in columns (1)–(4) of Table 2.[14] Column (1) presents a specification without firm fundamentals, while column (2), which we take as our benchmark, adds the main set of controls. Columns (3) and (4) restrict the sample to Census years only: column (3) mirrors the benchmark specification in column (2), and column (4) further incorporates the additional controls. Across all specifications, the debt ratio is positively and significantly associated with the probability of exit. The estimated coefficient is largely unchanged across specifications and increases only marginally when additional fundamentals are included.[15] Results are not only statistically significant but also economically relevant: in the benchmark specification, a one standard deviation increase in the debt ratio raises the propensity to exit by 1.02 percentage points, which corresponds to more than a 32 percent increase relative to the baseline exit rate.

Size and age interactions. To investigate whether the relationship between debt and exit is heterogeneous across firms, we extend the baseline specification by interacting the debt ratio with measures of firm size and age:

$$Exit_{ist} \;=\; \alpha\, DebtRatio_{is,t-1} \;+\; \beta\, DebtRatio_{is,t-1} \times Z_{is,t-1} \;+\; X'_{is,t-1}\,\omega \;+\; \delta_{st} \;+\; \varepsilon_{ist}, \tag{4.2}$$

where $Z_{is,t-1}$ denotes either $Large_{is,t-1}$, a dummy equal to one if firm i's sales are above the sector–year median, or $Young_{is,t-1}$, defined above. Results, reported in columns (5) and (6), show no statistically significant heterogeneity by firm size or age, although the point estimates suggest that the debt–exit relationship is somewhat stronger for smaller and younger firms.

[14] Table A.2 shows that the results are robust to using the debt-to-assets ratio instead of net short-term debt to sales.

[15] The comparison between columns (2) and (3) shows that results for Census years are in line with the full sample, which in turn justifies using column (4) to establish that additional controls add little beyond the benchmark specification.

Table 2: Propensity to Exit and Debt

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
debt ratio	0.010***	0.010***	0.009***	0.011***	0.011***	0.010***	0.009***
	(0.001)	(0.001)	(0.002)	(0.003)	(0.002)	(0.001)	(0.001)
debt ratio × large					−0.001		
					(0.001)		
debt ratio × young						0.006	
						(0.004)	
debt ratio × crisis							0.007*
							(0.003)
Observations	119K	119K	19K	15.5K	119K	119K	119K
R^2	0.007	0.008	0.008	0.010	0.008	0.008	0.009
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Fundamentals							
Main	No	Yes	Yes	Yes	Yes	Yes	Yes
Additional	No	No	No	Yes	No	No	No
Sample	Full	Full	Census	Census	Full	Full	Full

Notes: The dependent variable is a firm-level exit indicator equal to one if the firm exits the market in year t. The debt ratio is defined as net short-term debt to sales in year $t-1$. Main firm controls include a young firm indicator (five years or younger), firm size fixed effects (sales quintiles at the sector–year level), sales growth, and (log) capital stock; additional controls further account for (log) employment and (log) labor productivity. *Crisis* is an indicator equal to one if year t falls within an NBER recession. All regressions use the full sample, except columns (3) and (4), which use the Census sample. All regressions include industry–year fixed effects. Standard errors (in parentheses) are clustered at the industry–year level. Statistical significance at the 10, 5, and 1 percent levels is denoted by *, **, and ***, respectively.

Over the business cycle. We finally examine whether the relationship between debt and exit varies systematically over the business cycle. The specification is

$$\begin{aligned} Exit_{ist} \;=\; & \alpha\, DebtRatio_{is,t-1} \;+\; \gamma\, DebtRatio_{is,t-1} \times Crisis_t \\ & +\; X'_{is,t-1}\,\omega \;+\; X'_{is,t-1}\,\varphi \times Crisis_t \;+\; \delta_{st} \;+\; \varepsilon_{ist}, \end{aligned} \tag{4.3}$$

where $Crisis_t$ is a dummy equal to one if year t falls within an NBER recession. Controls are also interacted with the crisis indicator to allow their effects to vary over the cycle. The results, reported in column (7) of Table 2, show that the correlation between debt and exit strengthens significantly during recessions. The effect is not only statistically significant, but also economically meaningful: a one standard deviation increase in the debt ratio raises the probability of exit by approximately 0.92 percentage points (28.9 percent) in normal times, but by 1.63 percentage points (51.5 percent) in recessions.

4.3 Evidence from Italy

To complement the U.S. results, we exploit firm-level data from Italy collected from the Orbis database produced by Bureau van Dijk. Orbis provides annual balance sheet and income statement information for close to the universe of firms, allowing us to construct measures of financial health that parallel those used in the U.S. analysis.[16] Extending the analysis to Italy serves two purposes. First, it allows us to

[16] Appendix B describes the Orbis data, sample construction, and validation in detail.

Table 3: Propensity to Exit and Debt: U.S. vs. Italy

	Baseline			Over the cycle		
	(1)	(2)	(3)	(4)	(5)	(6)
	U.S.	Italy	Italy (QFR)	U.S.	Italy	Italy (QFR)
debt ratio	0.010***	0.032***	0.034***	0.009***	0.027***	0.029***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
debt ratio × crisis				0.007*	0.012***	0.012***
				(0.003)	(0.002)	(0.002)
Observations	119K	1.53M	1.3M	119K	1.53M	1.3M
R^2	0.008	0.031	0.030	0.009	0.031	0.030
Industry-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm Fundamentals	Yes	Yes	Yes	Yes	Yes	Yes

Notes: The dependent variable is a firm-level exit indicator equal to one if the firm exits the market in year t. The debt ratio is defined as net short-term debt to sales in year $t-1$. All specifications include main firm controls (young firm indicator, sales growth, log capital stock, and firm-size fixed effects) and industry–year fixed effects. *Crisis* is an indicator for NBER recessions (U.S.) or the euro-area crisis, 2008–2012 (Italy). The Italy (QFR) subsample restricts to firms with total assets above €311,000 in constant 2010 euros, the Italian equivalent of the QFR's $250,000 in 2000 asset cutoff for manufacturing firms. Standard errors (in parentheses) are clustered at the industry–year level. Statistical significance at the 10, 5, and 1 percent levels is denoted by *, **, and ***, respectively.

evaluate whether the debt–exit relationship generalizes beyond the U.S. context. Second, because Orbis covers the full firm size distribution, it enables us to assess whether the QFR's underrepresentation of small firms materially affects the estimated magnitudes.

A key advantage of Orbis over the QFR is that it covers the full firm-size distribution, including very small firms that fall below the QFR's asset threshold. We use this coverage to ask whether the QFR's underrepresentation of small firms distorts the estimated debt–exit relationship. We re-estimate specifications (4.1) and (4.3) on two samples: the full population of Italian manufacturing firms, and a QFR-comparable subsample that retains only firms above the Italian equivalent of the QFR's $250,000 asset cutoff.[17] Comparing the two isolates how much the exclusion of small firms matters for the estimates.

Table 3 reports the results. Two findings stand out. First, the positive and statistically significant association between debt and exit extends to Italy, confirming the external validity of the U.S. results. The estimated coefficients are roughly three times larger than in the United States. Through the lens of the model, this suggests a higher degree of financial frictions in Italy, consistent with well-documented differences in credit market development across the two economies. The amplification of the debt–exit link during crises, documented for the U.S. in specification (4.3), is also present: the interaction with the euro-area crisis indicator (2008–2012) is positive and highly significant.

Second, restricting the Italian sample to QFR-comparable firms has only a modest effect on the estimates. The baseline coefficient increases slightly from 0.032 to 0.034, and the crisis interaction is virtually unchanged. This result is reassuring for interpreting the QFR-based estimates: the exclusion of very small firms from the U.S. sample does not appear to materially distort the estimated relationship between debt and exit.

[17] The threshold is €311,000 in constant 2010 euros. A firm enters the QFR-comparable subsample in the first year it crosses this threshold and remains thereafter, mirroring the QFR's sampling design, under which firms stay in the survey regardless of subsequent asset fluctuations.

To summarize, the key finding of this section is that debt is a strong predictor of firm exit, even after controlling for non-financial fundamentals. This relationship is robust across specifications, holds in both the U.S. and Italy, and strengthens during recessions. The baseline coefficient of 0.010 from specification (4.1) is the empirical counterpart of the marginal propensity to exit with debt characterized in Section 3. As shown there, this moment is tightly linked to the degree of financial frictions in the model and will serve as a key calibration target in the next section.

5 Quantitative Analysis

In this section, we bring the model to the data. We calibrate the stationary equilibrium to match key features of the U.S. economy, including firm exit, growth, and the sensitivity of exit to leverage documented in Section 4. We then use the calibrated model to quantify the welfare costs of financial frictions operating through firm exit. We also study the economy's response to a recessionary shock, validating the model's ability to generate an increase in the marginal propensity to exit with debt during downturns, consistent with the empirical evidence. Finally, we evaluate the effectiveness of government interventions, including policies in the spirit of the Paycheck Protection Program.

5.1 Calibration

Our calibration strategy is closest to Karahan, Pugsley, and Şahin (2024), who calibrate a firm dynamics model to match age-specific exit rates, startup size, and firm growth profiles from the Business Dynamics Statistics (BDS), building on the empirical findings of Pugsley, Sedláček, and Sterk (2021). The key difference between their model and ours is the presence of financial frictions. Absent financial frictions, i.e., with $\kappa = 0$, our economy closely resembles theirs: default has no productivity consequences, so debt is irrelevant for exit and the allocation is efficient. We adopt their set of targeted moments and augment it with the marginal propensity to exit with debt. As shown in Section 3, this moment is tightly linked to the degree of financial frictions and is therefore key for identifying κ.

Functional forms. The quantitative model requires specifying the productivity process and the distribution of fixed operating costs. For productivity, we follow Pugsley, Sedláček, and Sterk (2021) and decompose idiosyncratic firm productivity into a permanent and a transitory component. Total idiosyncratic productivity is $a_{jt} = a_j^p \cdot a_{jt}^{tr}$, where a_j^p is a permanent type drawn at entry and a_{jt}^{tr} evolves over a firm's lifetime.[18] The permanent component takes three values and is drawn at entry from a discretized log-normal distribution with standard deviation σ_a^p and unit mean. The transitory component follows an AR(1) process in logs:

$$\log a_{jt}^{tr} = \rho_a \log a_{j,t-1}^{tr} + \sigma_a u_{jt} \quad \text{with} \quad u_{jt} \sim N(0,1),$$

[18] As shown by Pugsley, Sedláček, and Sterk (2021), allowing for permanent productivity differences across firms is essential for matching the observed patterns of firm growth. Models that rely solely on persistent transitory shocks overpredict the dispersion in growth rates among older firms. Permanent heterogeneity, by contrast, allows the model to jointly capture the high exit and fast growth of young firms alongside the flattening of these profiles with age.

where ρ_a governs persistence and σ_a is the standard deviation of innovations. Entrants draw their initial transitory productivity from a log-normal distribution with mean μ_a^{ent} and standard deviation σ_a^{ent}.

For fixed operating costs, we assume that f_{jt} is drawn each period from a Pareto distribution with scale parameter $\underline{f}$ and shape parameter χ_f. The Pareto specification generates a thick right tail, so that large cost realizations occur with non-negligible probability even though the typical draw is moderate. This feature is important for producing realistic exit rates that decline with firm age but remain positive even for old, productive firms.

Externally calibrated parameters. We calibrate the model at an annual frequency. Two parameters are set externally. The discount factor is $\beta = 0.96$, corresponding to a risk-free rate of approximately 4%. The elasticity of substitution across varieties is $\epsilon = 3$, a standard value in the firm dynamics literature.

Internally calibrated parameters. The remaining nine parameters are calibrated by matching moments of the stationary equilibrium to their data counterparts. Table 4 reports the calibrated values. These parameters govern the distribution and dynamics of firm productivity, σ_a^p, ρ_a, and σ_a, the entrant productivity distribution, μ_a^{ent} and σ_a^{ent}, the marginal entry cost, $\bar{f}_e$, the distribution of fixed operating costs, $\underline{f}$ and χ_f, and financial frictions, κ.[19]

Table 4 reports the nine targeted moments together with the model's fit. The first eight are firm dynamics moments from Karahan, Pugsley, and Şahin (2024) and Pugsley, Sedláček, and Sterk (2021), computed using the Census micro data. These moments capture exit rates by firm age, computed as firm-weighted averages within five-year bins (ages 1–5, 6–10, 11–15, and 16–19), average startup size measured as employment at age one, and employment growth rates of surviving firms at age three for three initial size groups: small (1–49 employees), medium (50–249), and large (250 or more). The ninth moment is the marginal-propensity-of-exit coefficient from our baseline regression specification in Table 2. In the model, the analogous coefficient is obtained by regressing exit probability on debt-to-sales with productivity fixed effects. The model matches all nine moments closely.

Identification. While all nine parameters are jointly calibrated, the mapping from parameters to moments has an intuitive structure that aids identification.

The four binned exit rates discipline the productivity process and the distribution of operating costs. The average exit rate in the economy helps pin down the operating cost parameters, $\underline{f}$ and χ_f: a higher scale $\underline{f}$ or a lower shape χ_f, corresponding to a thicker tail, raises exit rates across all ages. The slope of exit rates across age bins is governed by the productivity process. Higher persistence, ρ_a, slows the rate at which firms are selected, flattening the age profile of exit. A larger innovation standard deviation, σ_a, increases the dispersion of productivity among incumbents and steepens the decline in exit rates as low-productivity firms are weeded out. The relation between the operating cost parameters also helps shape the exit-by-age profile. By increasing the tail of the Pareto distribution

[19] The calibration procedure uses f_e which governs the distribution of entry costs among potential entrants. We report the equilibrium value of the marginal entry cost, $\bar{f}_e$, as it has a clearer economic interpretation.

Table 4: Parameters and Targeted Moments

Panel (a): Externally Calibrated Parameters					
Parameter	Description	Value	Source		
β	Discount factor	0.96	Risk-free rate $\approx 4\%$		
ϵ	Elasticity of substitution	3	Standard value		
Panel (b): Internally Calibrated Parameters					
Parameter	Description	Value	Moment	Data	Model
$\underline{f}$	Operating cost scale (Pareto)	0.50	Exit rate, ages 1–5	0.13	0.13
χ_f	Operating cost shape (Pareto)	0.95	Exit rate, ages 6–10	0.08	0.08
σ_a	Transitory productivity innov. s.d.	0.23	Exit rate, ages 11–15	0.06	0.06
ρ_a	Transitory productivity persistence	0.96	Exit rate, ages 16–19	0.05	0.05
σ_a^p	Permanent productivity s.d.	0.22	Growth, small firms	1.08	1.08
μ_a^{ent}	Entrant mean log productivity	−0.34	Growth, medium firms	1.01	1.01
$\sigma_{\underline{a}}^{ent}$	Entrant productivity s.d.	0.42	Growth, large firms	0.97	0.97
$\bar{f}_e$	Marginal entry cost	35.66	Startup size (employees)	6.10	6.10
κ	Financial friction	0.25	debt–exit coefficient	0.01	0.01

Notes: The table presents the calibrated parameters and the moments the calibration targets. Growth rates are employment-weighted growth ratios for surviving three-year-old firms, by initial size group: small (1–49 employees), medium (50–249), and large (250+). Data on exit rates, startup size, and growth rates are taken from Karahan, Pugsley, and Şahin (2024) and Pugsley, Sedláček, and Sterk (2021). The debt–exit coefficient is taken from Table 2.

(lower χ_f) and simultaneously lowering the scale parameter ($\underline{f}$), the model can hit the same average exit rate in the economy while having a flatter exit-by-age profile.

Average startup size is primarily pinned down by the marginal entry cost, $\bar{f}_e$: a lower cost raises the entry rate, increasing the number of operating firms and lowering their individual employment levels. This acts as the main lever for matching the 6.1 employees per entrant observed in the data. The growth rates by initial size help separate the permanent and transitory components of productivity. Greater permanent heterogeneity, σ_a^p, generates dispersion in the initial size of entrants. Conditional on starting small, firms with low permanent productivity have limited growth prospects, while those with high permanent productivity grow rapidly. The relative magnitude of σ_a^p and σ_a^{ent} determines how much of the variation in initial size reflects permanent types rather than transitory draws, thereby shaping conditional growth rates. In particular, the declining growth profile across initial size groups, from 8.1% for small firms to −3.0% for large firms, disciplines the degree of mean reversion in the transitory component and the role of permanent heterogeneity.

The entrant productivity distribution, governed by μ_a^{ent} and σ_a^{ent}, further shapes the initial conditions of entering cohorts. A lower mean, μ_a^{ent}, implies that entrants start with below-average productivity, consistent with the evidence in Pugsley, Sedláček, and Sterk (2021) that the fast growth of surviving young firms reflects selection, while a larger dispersion, σ_a^{ent}, increases heterogeneity in initial draws and affects both startup size and subsequent growth.

Let us now turn to the identification of the key parameter of our model, κ. This parameter governs the degree of financial frictions and is primarily identified by the regression coefficient of exit on debt-to-sales. When $\kappa = 0$, so that there are no financial frictions, Proposition 2 established that debt is irrelevant for exit and the regression coefficient equals zero. Proposition 3 highlighted that in the simplified version of our model, a higher κ increases the marginal propensity to exit with debt. As κ

Figure 1: Identification of the degree of financial friction

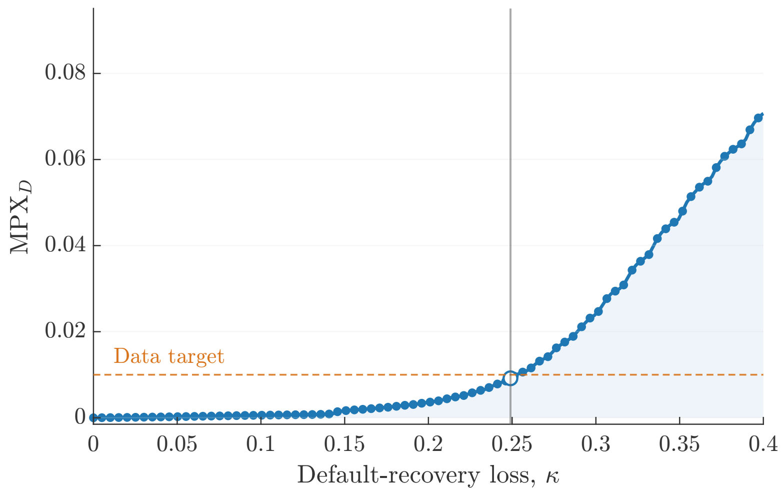

Notes: This figure shows how the marginal propensity to exit with debt varies with κ in the stationary distribution, holding other parameters fixed at their calibrated values. The strong positive relationship sheds light on how κ is identified.

rises, highly leveraged firms are more likely to exit because default triggers a productivity loss that can make operating unprofitable under new ownership. The calibrated value, $\kappa = 0.25$, implies that when a firm defaults, new owners operate it at roughly 75% of its original productivity. Figure 1 presents comparative statics with respect to κ, plotting the model-implied regression coefficient as a function of κ, holding the other structural parameters fixed at their calibrated values. The increasing relationship between financial frictions and the marginal propensity to exit with debt shows that the result of Proposition 3 extends to our quantitative model.[20]

The calibration of κ hinges on the estimated regression coefficient (4.1). One may worry that despite controlling for observable firm-level characteristics in the empirical specification, there could be unobservable characteristics that affect the relationship between debt and exit. Ultimately, the parameter κ affects the debt pricing schedule as well as the borrowing limit through the recovery rate following default—the net-present-value of the firm following the productivity loss over the face value of the loan. Therefore, an alternative strategy for calibrating κ would be to match the observed recovery rate in the data. Over the period 1982–2025, the average recovery rate on unsecured bank loans to businesses in the U.S. is 44% (Moody's Investors Service, 2026). While we do not target this moment in our calibration, the model-implied recovery rate is remarkably close, equaling 43.2%. The tight fit of this untargeted moment implies that our quantitative analysis is robust to this alternative calibration strategy.

5.2 Features of the Calibrated Economy

We now describe the key features of the calibrated stationary equilibrium.

[20]While the figure shows that this relationship holds when we hold other parameters fixed at their calibrated values, Figure A.1 in the Appendix shows that this pattern continues to hold when we allow other model parameters to vary.

Figure 2: Ergodic Distribution of Firms

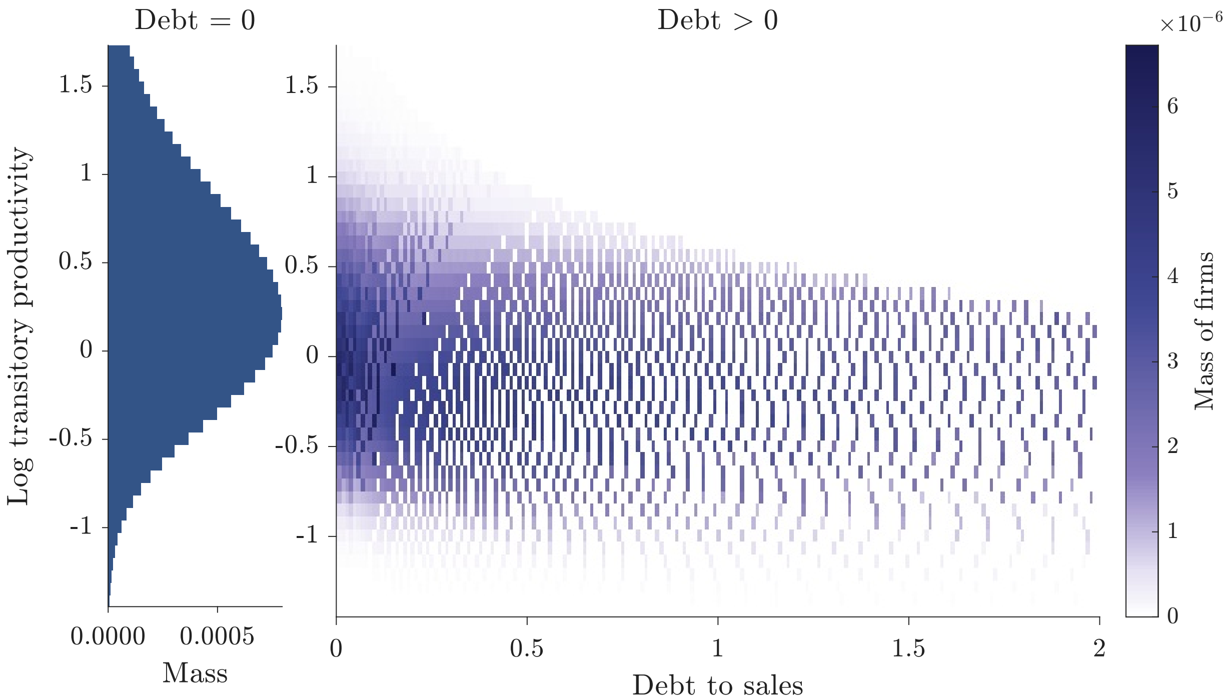

Notes: This figure displays the stationary distribution of firms over transitory log productivity and the ratio of initial debt obligations to sales. This figure considers only firms with medium permanent productivity, but similar patterns arise for firms with low and high permanent productivity. The left panel shows firms with zero debt; the right panel shows the distribution of firms with positive debt.

Ergodic distribution. Figure 2 displays the ergodic distribution of firms over transitory productivity and their debt-to-sales ratio. The figure shows this distribution for firms with medium permanent productivity, which account for 69% of operating firms in the economy.[21] The left panel shows the distribution of firms with zero debt. There is a mass at zero debt because firms enter the economy with no debt and a substantial fraction of firms has positive cash flows. The right panel shows the distribution of firms with positive debt. Across all productivity levels, 22.6% of firms have positive amounts of debt. This number is remarkably close to the share of firms with positive net short-term debt position in the QFR dataset, which is 24.4%. The indebted firms are concentrated around the mean of the productivity distribution. High-productivity firms finance their operations out of sales and borrow little, while the mass of low-productivity firms is small because of survival selection.

Life-cycle dynamics. Figure 3 reports the life-cycle dynamics of firms. As firms age, mean reversion and selection raise their productivity: a 10-year-old operating firm is about 50 log points more productive than the average entrant, and its average size grows from 6.1 workers at entry to about 20 workers by age 10. Since productivity is back-loaded, young firms are willing to devote a larger share of their sales to fixed operating costs. The share of workers not engaged in direct production falls from 25% for entrants to less than 10% for mature firms.

Debt, default, and exit also evolve over the life cycle. Entrants begin without debt and are less

[21] Figure A.2 in the Appendix corresponds to low- and high-permanent productivity levels.

Figure 3: Equilibrium Age Profiles

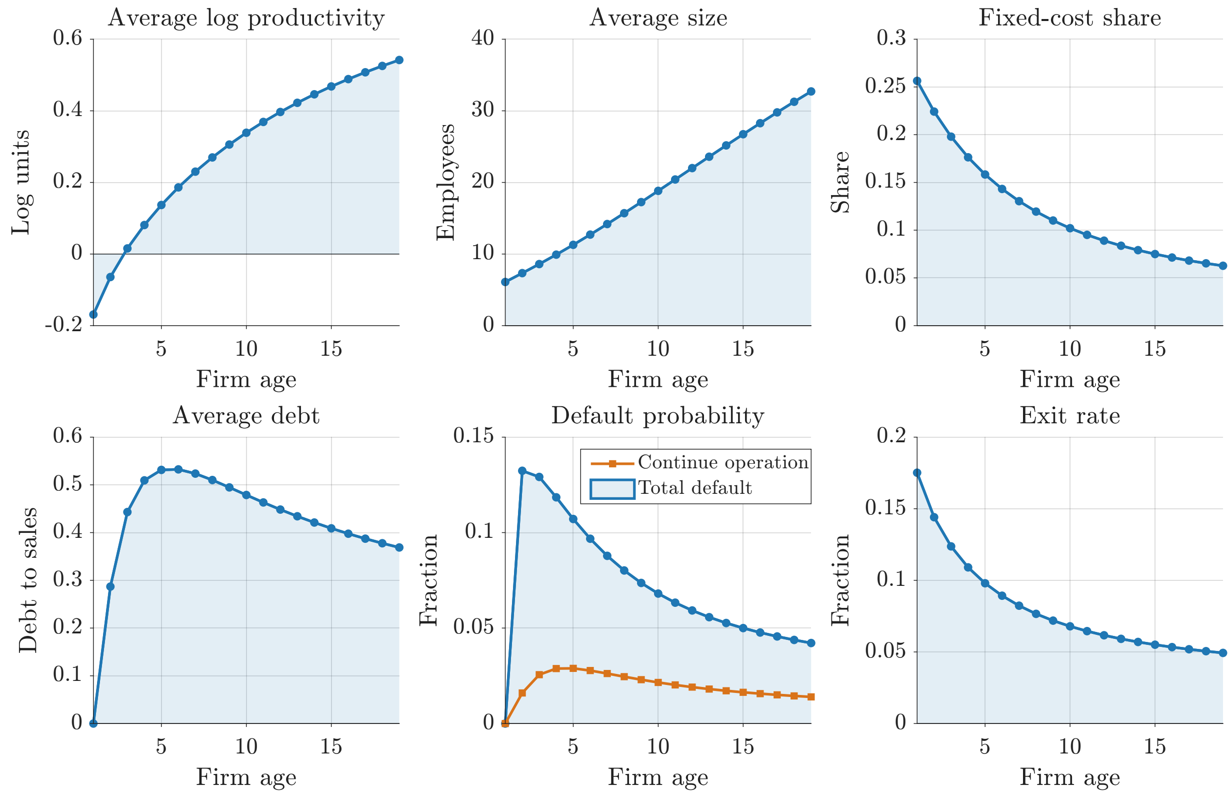

Notes: This figure displays average values of key firm-level variables by firm age in the stationary equilibrium. Exit rate is the fraction of firms that exit prior to the next period. Average size is measured in employees. The fixed-cost share indicates the share of workers not engaged in direct production. Default probability is the fraction of firms that default (but may continue under new ownership). Log productivity combines permanent and transitory components.

productive than average, so they borrow to finance operations, and the debt-to-sales ratio rises over the first five years. As productivity climbs, the need to borrow declines and the ratio falls with age. Default and exit follow. Entrants have zero default probability by construction, but can still exit when a high fixed-cost draw exceeds their continuation value; the default probability rises to about 13% by age two. Most defaulting firms exit, because financial intermediaries rarely find it profitable to continue operating them: the share of firms that default and continue under new ownership peaks at only 3% at age four. The exit rate declines throughout the life cycle, from about 18% for entrants to about 5% for firms older than fifteen.

Exit and default probabilities. It is instructive to delve deeper into the exit and default probabilities over the state space. Figure 4 shows the exit (left panel) and default (right panel) probabilities for a firm with medium permanent productivity over the state space.[22] Absent financial frictions, exit probability would be independent of debt. In the presence of financial frictions, i.e., when κ is strictly positive, more indebted firms are more likely to exit. This is the pattern observed in the left panel of

[22] Figures A.3 and A.4 in the Appendix report the corresponding patterns for firms with low and high permanent productivity.

Figure 4: Exit and default patterns

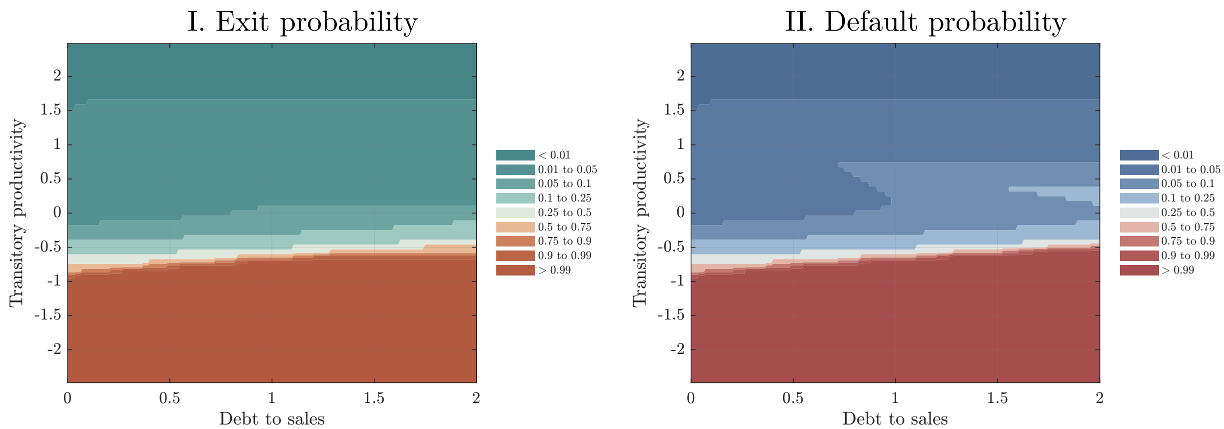

Notes: This figure presents the exit (left panel) and default (right panel) probabilities as a function of the start-of-period transitory productivity (y-axis) and debt-to-sales (x-axis) levels, for a medium permanent productivity firm. The higher is the firm productivity, the lower is its default and exit probability. The higher is its debt-to-sales ratio, the higher are these probabilities.

Figure 4. Exit rates are increasing in debt levels and decreasing in the level of productivity. The right panel reveals a similar pattern for default probabilities.[23]

5.3 The costs of financial frictions

To quantify the costs of financial frictions operating through firm exit, we compare the calibrated economy ($\kappa = 0.25$) to a frictionless benchmark in which $\kappa = 0$. Proposition 2 established that this benchmark coincides with the efficient allocation: default entails no productivity loss, so debt is irrelevant for exit and all firm turnover is driven by fundamentals. Table 5 reports the comparison. The welfare gain from eliminating financial frictions is 3.6% in consumption-equivalent terms, accounting for the full transition path from the baseline to the frictionless steady state.

The steady-state comparison reveals how financial frictions reshape the firm distribution. In the baseline, 9.3% of firms exit each period. Removing financial frictions reduces exit rate to 5.0%, a decline of 46%. The additional incumbents raise the mass of operating firms by 55%, from 0.034 to 0.053. As more firms share the same aggregate labor supply, individual firms become smaller: average employment falls from 30.3 to 19.0 workers. Steady state consumption rises by 8.3%. The source of this gain is straightforward: financial frictions force out firms whose present discounted social value is positive but whose borrowing capacity falls short of their current obligations. Eliminating these frictions allows such firms to survive, raising the measure of available varieties and improving resource allocation along the extensive margin.

Figure 5 shows how exit probabilities change in the first period of the transition from the baseline to the frictionless economy, holding the distribution of firms fixed at the baseline ergodic distribution. This exercise isolates the effect of removing financial frictions for a given population of firms. Each panel plots the difference in exit probabilities as a function of transitory productivity and the debt-

[23] While default probabilities are monotone in the productivity level for a given level of debt, they can decline for a given level of debt-to-sales ratio. This non-monotonic pattern is driven by the productivity drop following default.

Table 5: Costs of Financial Frictions

	Baseline	No financial frictions	Change
Mass of firms	0.034	0.053	+54.9%
Exit rate	9.3%	5.0%	−46.2%
Average size (employees)	30.3	19.0	−37.2%
Steady-state consumption	0.387	0.419	+8.3%
Welfare gain (CE, with transition)			+3.6%

Notes: The table compares steady-state outcomes under the calibrated and the frictionless benchmark ($\kappa = 0$). The welfare gain is the consumption-equivalent variation, computed using the full transition path from the baseline to the $\kappa = 0$ steady state.

to-sales ratio for low, medium, and high permanent productivity, respectively. The reduction in exit is concentrated among firms with moderate-to-low productivity and high debt burdens. High-productivity firms rarely exit regardless of their financial position, because their borrowing capacity is large enough to cover even adverse operating cost realizations. Very low-productivity firms exit under both regimes, because their fundamentals do not justify continued operation. The firms that financial frictions push out are those in between: productive enough to survive absent debt overhang, but too leveraged to refinance when hit by a large fixed operating cost shock.

Removing financial frictions raises real wages in the economy, which compresses profits for all firms. This general equilibrium force partially offsets the direct effect of eliminating frictions. Some marginal low-debt firms that were viable in the baseline are pushed below the exit threshold in the frictionless economy, accounting for the small negative region in Figure 5. Higher wages also reduce the value of entry, lowering the equilibrium mass of entrants. Despite these offsetting forces, the net effect remains large.

The welfare costs reported in Table 5 reflect two distinct channels through which financial frictions reduce efficiency. The first is the exit channel: financial frictions prevent viable firms from refinancing, pushing them out of the market. The second is a direct productivity loss channel: firms that default but continue operating under new ownership produce at a fraction $(1 - \kappa)$ of their original productivity. Even when these firms survive, the accumulated productivity losses reduce their output and lower aggregate consumption. Because transitory productivity follows an AR(1) process, the scars of past defaults partially fade over time through mean reversion. To isolate the contribution of the productivity loss channel, we take the baseline ergodic distribution of operating firms and track the cumulative losses to productivity attributable to past defaults. We then compute the counterfactual level of aggregate consumption that would obtain if each firm's transitory productivity were restored to its pre-default path.[24] Restoring defaulters' productivity raises aggregate consumption by 2.2%, accounting for roughly one quarter of the 8.3% steady-state consumption gap. The remaining three quarters are attributable to the exit channel. The modest size of the productivity loss channel reflects the fact that few defaulting firms continue operating: only 1.6% of firms default and survive under new ownership each period, and just 5.9% of operating firms carry productivity scars from a past default.

[24]This calculation holds firm decisions, prices, and the set of operating firms fixed. It therefore isolates the mechanical output loss from accumulated productivity scarring, without confounding it with changes in exit behavior or general equilibrium responses.

Figure 5: Excess Firm Exit

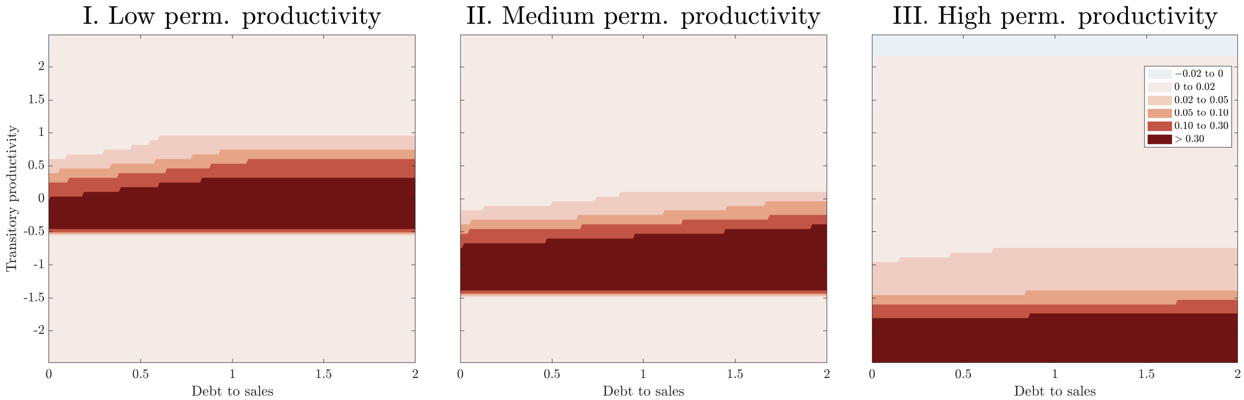

Notes: This figure displays the difference in exit probabilities between the baseline economy and the first period of the transition towards $\kappa = 0$, holding the distribution of firms fixed at the baseline ergodic distribution. Panels I, II, and III correspond to low, medium, and high permanent productivity, respectively. For each panel, the horizontal axis measures the debt-to-sales ratio and the vertical axis measures transitory productivity.

5.4 Recessions

The analysis so far has focused on steady-state comparisons. We now study the economy's response to two types of recessionary shocks: an adverse aggregate productivity shock and a financial crisis. The productivity shock captures a standard business cycle downturn driven by fundamentals. The financial crisis captures an episode in which credit markets deteriorate. The two experiments deliver a sharp contrast. Productivity shocks barely move exit rates. Financial shocks tighten firms' borrowing capacity, generating large exit responses with sizable welfare consequences.

Productivity shock. We model a standard recession as a one-time unexpected decline in aggregate productivity Z_t. On impact, Z_t drops by 5 log points and reverts to its steady-state value according to an AR(1) process with persistence $\rho_Z = 0.5$, so that its half-life is one year. Figure 6 reports the impulse response of the baseline economy.

Exit responds only modestly to the productivity shock. On impact, the exit rate rises by 27 basis points. The increase is short-lived: within 3 periods, exit falls *below* its steady-state level. Two forces drive this reversal. First, the initial spike in exit is cleansing. Firms pushed out by the shock are disproportionately low-productivity, so the surviving population is more productive on average. Fewer incumbents draw operating costs that exceed their continuation value, and exit declines. Second, the drop in entry reduces the inflow of young, fragile firms in subsequent periods, further suppressing exit during the recovery. Entry itself falls by approximately 40 basis points on impact, as lower aggregate productivity reduces the expected present value of a new firm. The cumulative net exit over the first five years equals 60 basis points.

Table 6 studies the welfare costs of the recession. Households in the economy are indifferent between going through the recession and suffering a 1-year consumption decline of 11.17%. To understand the role financial frictions play in the recession, we compute the magnitude of *excess exit* both in the

Figure 6: Impulse Response to a Productivity Shock

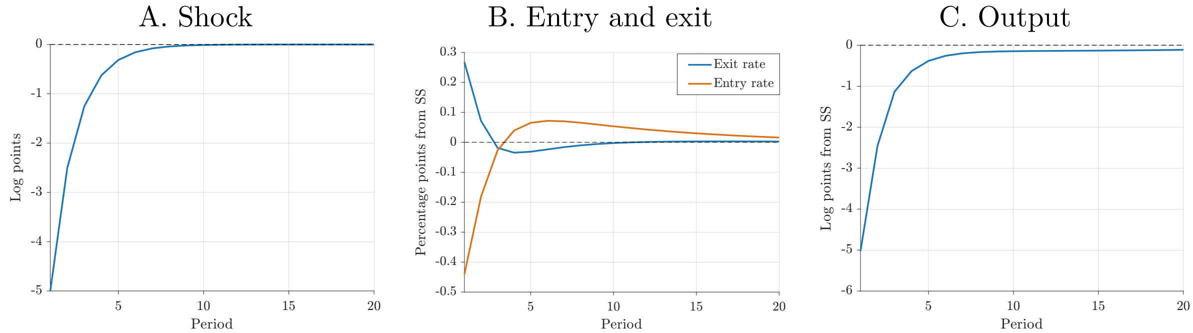

Notes: This figure displays the impulse response to a one-time 5 log points decline in aggregate productivity Z_t with persistence $\rho_Z = 0.5$. Panel A plots aggregate TFP in log deviations from steady state. Panel B plots entry and exit rates in percentage-point deviations from their respective steady-state values. Panel C plots aggregate output in log deviations from steady state.

steady state and at the onset of the recession. For each firm, excess exit is defined as the difference between the equilibrium exit rate and the exit rate a planner would choose in the first period, assuming the firm is subject to financial frictions from the next period on. In the steady state, the excess exit rate equals 4.65pp. In the recession, the excess exit rate increases by only 9 basis points to 4.74pp. The small changes in both the exit rate and excess exit rate indicate that financial frictions do not amplify a productivity-driven recession through the exit margin.

The acyclicality of the exit rate is consistent with the data. Exit rates in the United States are stable across most postwar recessions, rising by less than one percentage point even in severe downturns (see Lee and Mukoyama, 2015; Ayres and Raveendranathan, 2023). The Great Recession stands out as the exception. We turn to it next.

Financial crisis. We model a financial crisis as an unexpected increase in the financial friction parameter κ_t. In our framework, κ governs recovery values, borrowing capacity, and the sensitivity of exit to debt. A shock to κ therefore captures the key features of a financial crisis: collapsing recovery values, tighter credit markets, and a stronger link between debt and firm survival. We assume that the persistence of the shock to κ equals 0.5, as we did for the aggregate productivity shock, and calibrate the size of the shock to match the rise in the exit rate during the Great Recession, 1.67 percentage points.[25]

The financial crisis operates through a fundamentally different channel than the productivity shock. Under a productivity shock, exit rises because lower TFP reduces the present value of continuation for marginal firms. The firms pushed out are those with the weakest fundamentals. Under a financial crisis, exit rises because borrowing capacity contracts. The firms pushed out are those identified in Figure 5: moderate-to-low productivity, high debt burdens, productive enough to survive under normal financial conditions but too indebted to weather the decline in borrowing capacity. As a result, exit

[25] We measure exit rates from the BDS as firm deaths divided by the average of current and prior-year firm counts, and compute deviations from a linear trend over 1979–2019. The 1.67 percentage point value corresponds to the peak deviation in 2009.

Table 6: Welfare Costs of Recessions

	Welfare cost (1-yr CE %)	Δ Exit rate (year 1)	Δ Excess exit rate (year 1)
Productivity shock	11.17	0.27	0.09
Financial crisis	3.61	1.67	1.38

Notes: Welfare costs are measured as the 1-year percentage decline in consumption that yields the same present discounted utility loss as the shock. Δ exit rate and Δ excess exit rate report the change, in percentage points, between the steady state and first period exit and excess exit rates, respectively. The productivity shock is a 5 log points decline in Z_t with persistence $\rho_Z = 0.5$. The financial crisis is an increase in κ_t from 0.25 to 0.30 with persistence $\rho_\kappa = 0.5$.

is less cleansing than in a standard recession. Firms that would generate positive social value are destroyed because they cannot raise the credit to continue operating.

The entry and output dynamics reinforce this distinction. Panel B of Figure 7 shows that entry initially dips but then rises *above* its steady-state level. The wave of exit reduces the mass of operating firms, which raises profits for survivors and makes entry more attractive. This stands in contrast to the productivity shock, where entry falls sharply and remains depressed throughout most of the downturn.

Panel C reveals that the output response to a financial crisis is much more persistent. On impact, output rises slightly. This is because exiting firms release workers from fixed operating activities, and these workers are reallocated to direct production at the surviving firms. In the short run, this reallocation effect dominates the loss of variety. Over time, however, the destruction of firms cumulates. Output falls persistently, reaching 0.3 log points below steady state by period 15, even though κ_t has largely reverted by then. The source of this persistence is the slow process of rebuilding the stock of operating firms through entry. Unlike a productivity shock, where output recovers as TFP reverts, here the damage is long lasting. Overall, the welfare cost associated with the financial crisis is 3.61% in one-year consumption-equivalent terms, as presented in Table 6. Our aim is not to compare the welfare costs of a financial crisis to a 5 log point decline in productivity. Rather, we emphasize that the costs from excess exit are substantially larger in a financial crisis.[26]

In summary, in a standard productivity-driven recession, excess exit barely rises, so financial frictions add little to the downturn through the extensive margin. In a financial crisis, excess exit jumps and generates sizable, persistent welfare losses. Government intervention to support firm survival can mitigate these costs. We evaluate two such interventions next.

5.5 Policy

We use the calibrated model to evaluate two policies inspired by programs deployed during the COVID-19 pandemic: a grant in the spirit of the Paycheck Protection Program (PPP) and a government-guaranteed loan in the spirit of the Economic Injury Disaster Loan (EIDL) program. We study these programs in the steady state of the economy rather than during a financial crisis. This isolates the mechanics of each instrument and provides a clean benchmark for assessing their effectiveness.

[26] In principle, because firm fundamentals are unchanged, all additional exit induced by the financial crisis is inefficient. The "excess exit" measure in the final column of Table 6, however, is based on a one-period planner problem that takes future financial frictions as given. Since persistent frictions reduce the planner's value of keeping firms alive, the rise in measured excess exit need not equal the total rise in exit.

Figure 7: Impulse Response to a Financial Crisis

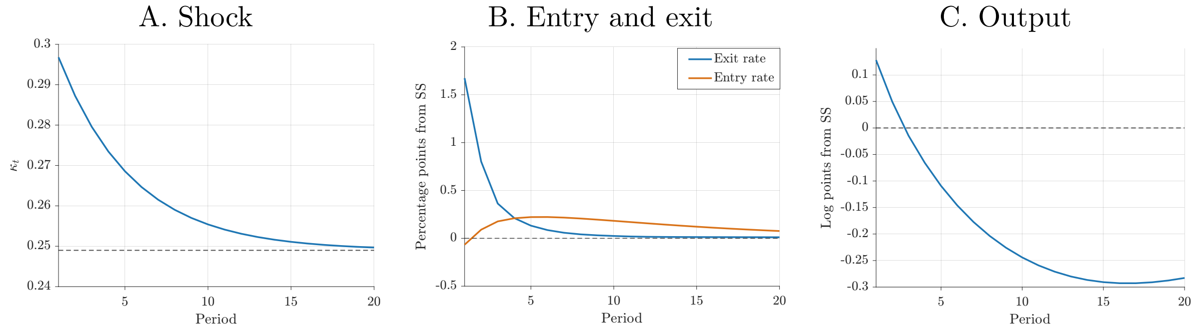

Notes: This figure displays the impulse response to a one-time increase in κ_t from 0.25 to 0.30 with persistence $\rho_\kappa = 0.5$. Panel A plots the path of κ_t. Panel B plots entry and exit rates in percentage-point deviations from their respective steady-state values. Panel C plots aggregate output in log deviations from steady state.

Grant program. The first policy is a one-period grant program. Each firm is eligible to receive a transfer of up to 20.8% of its prior-year wage bill, the equivalent of the 2.5-months-of-payroll formula used in the first round of the PPP. The transfer is only extended to firms with fewer than 500 workers and is further capped by the firm's cash-flow shortfall in the period, in line with the PPP's necessity certification: firms that can cover their operating costs without external support receive nothing. We assume that these transfers are financed by a lump-sum tax on households. We consider an unexpected implementation of the program for a single period, starting from the steady state of the decentralized economy, and study the transition dynamics of the economy back to the steady state and the resulting welfare implications.

The first row of Table 7 summarizes the welfare evaluation of the grant program. The program reduces the exit rate from 9.3% in the baseline economy to 7.6% on impact. The associated welfare gains equal 1.67% in 1-year consumption-equivalent terms. That is, households are indifferent between the implementation of the policy and receiving a one-time increase in consumption of 1.67%. The overall cost of the program is 0.75% of steady-state output.

To understand the welfare effects of the grant program, recall that our model is *statically efficient*, so all gains stem from reducing exit and default rates. Figure 8 sheds light on the heterogeneous effects of the policy across firms. The solid lines in Figure 8 present the exit rate in the steady state of the baseline economy for firms without debt along the productivity distribution. The dashed lines present the socially optimal exit rates: for each firm, this is the exit rate a planner would choose in the first period, assuming the firm is subject to financial frictions from the next period on. Due to financial frictions, throughout the state space, the efficient exit rates are below the decentralized-equilibrium ones. The solid lines with triangle markers present the exit rates in the first period following the program implementation. The grant reduces exit rates throughout the state space, bringing the economy closer to the dynamically-efficient allocation. While the aggregate welfare effects of the policy are positive, we note that the exit rates of some low-productivity firms fall below the socially-efficient ones.

Table 7: Welfare and Cost of Policy Interventions

	Gross cost (year 1, Y %)	Net cost (PDV, Y %)	Δ Exit rate (year 1, pp)	Welfare gain (1-yr CE %)
Grant	0.75	0.75	−1.75	1.67
Guaranteed loan	0.70	0.29	−0.93	1.23

Notes: Both programs transfer up to 20.8% of the firm's annual wage bill, capped by its cash-flow shortfall in the period. The guaranteed loan extends credit at $r_f + 2.5$ percentage points repaid over 30 years. Gross cost is the size of the fiscal transfer in year 1, as a fraction of steady-state output. Net cost is the present value of fiscal flows, taking into account firm loan repayments. The change in exit rate is relative to the steady state. The Welfare gain is the one-year consumption-equivalent variation.

Government-guaranteed loan program. The second policy we consider is a government guaranteed loan program with a thirty-year horizon. The government extends credit to firms at a fixed interest rate, 2.5 percentage points above the risk-free rate, repaid on a mortgage-style amortization schedule.[27] To simplify the comparison with the grant program, we assume that the loan size is capped at 20.8% of its prior-year wage bill, is only available for firms with prior-year employment below 500 workers, and cannot exceed the firm's operating cash-flow shortfall.[28] The government bears all credit losses, financed by lump-sum taxes. As in the grant-policy analysis, we consider an unexpected and temporary implementation for a single period.

The second row of Table 7 summarizes the guaranteed-loan program. The program reduces the exit rate from 9.3% in the baseline economy to 8.4% on impact, and delivers welfare gains of 1.23% in one-year consumption-equivalent terms, 26% below the grant. The loan is less effective at preventing exit even though it shares the same cap and eligibility conditions as the grant, and Figure 8 shows its smaller reduction in exit rates along the productivity distribution. The reason is debt overhang. Firms cannot selectively default; otherwise no firm would repay the guaranteed loans. The credit disbursed by the government therefore raises firms' default probabilities in future periods, lowering their endogenous borrowing limit already in the period of implementation, and leaving more firms exposed to exit.[29]

Despite this difference in effectiveness, the two programs disburse similar amounts: 0.7% of steady-state output for the loan against 0.75% for the grant. The near-equality reflects the imperfect targeting of both instruments. Each reaches roughly 30% of operating firms but prevents the exit of only 1–2%, suggesting that better targeting could deliver large welfare gains. The two programs differ sharply, however, in their net cost. The government recoups a large share of its lending. Over the thirty years following implementation, the government recovers 59% of the first-year disbursement, leaving a net present cost of 0.29% of steady-state output.

[27] EIDL loans carried a fixed interest rate of 3.75%. We measure the spread relative to the 10-year Treasury yield, which roughly matches the effective duration of amortizing, prepayable loans. With the 10-year yield around 1% over the main origination period in 2020, this implies a spread of roughly 250–300 basis points. We use 250 basis points as a benchmark.

[28] The EIDL program had a higher cap in practice, 50% of the firm's prior-year sales. In the final part of this section we do sensitivity analysis with respect to the size of the loan cap.

[29] This result is related to Crouzet and Tourre (2021), who study the effects of government-guaranteed loans on investment through debt overhang in a partial equilibrium environment without firm exit.

Figure 8: Exit Probability under Alternative Policy Scenarios

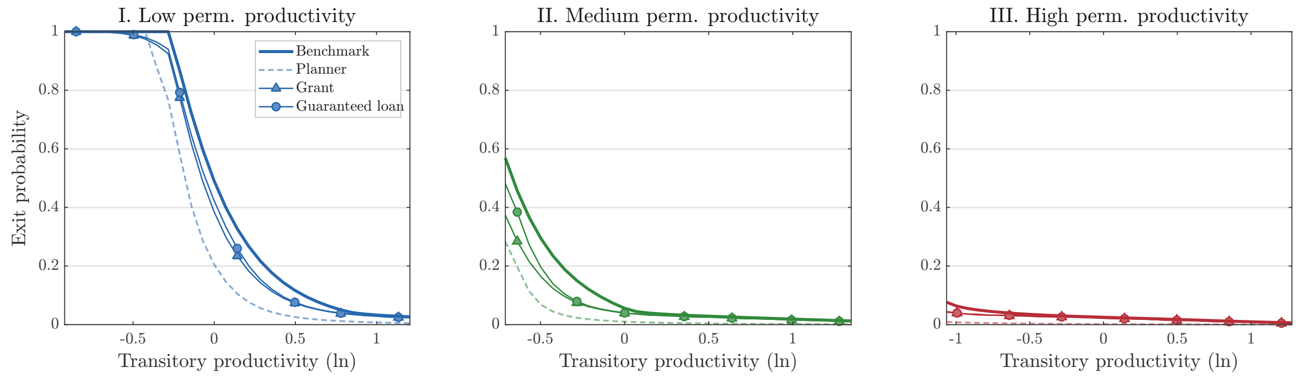

Notes: This figure displays the average exit probabilities as a function of a firm's productivity. Each panel compares four policies: the benchmark equilibrium, the planner's optimum, a targeted grant, and a guaranteed loan. Panels I, II, and III correspond to low, medium and high permanent productivity, respectively. The x-axis for each panel are chosen to cover 90% of operating firms in the corresponding productivity category, prior to the exit decision.

Welfare and program size. Figure 9 studies how the cost and welfare impact of the two programs vary with the maximum disbursement amount. The left panel of the figure presents the net costs of each program as a function of the prior-year wage-bill limit. These costs are monotonically increasing and concave. As the limit increases, the firm exit rate declines, and the aggregate disbursement amount increases. Recall that the disbursement amount per firm is limited by its cash-flow shortfall. Absent this restriction, doubling the limit of the program would result in more than double the cost—the disbursement amount per firm would double and the number of firms receiving a benefit would increase. With the cash flow shortfall limit in place, the disbursement amount per firm does not scale linearly with the benefit limit. Because of the decreasing density of the operating cost distribution, doubling the limit does not double the number of firms "saved" by the program. Overall, this results in a concave and monotonic relationship.

Figure 9: Program size sensitivity

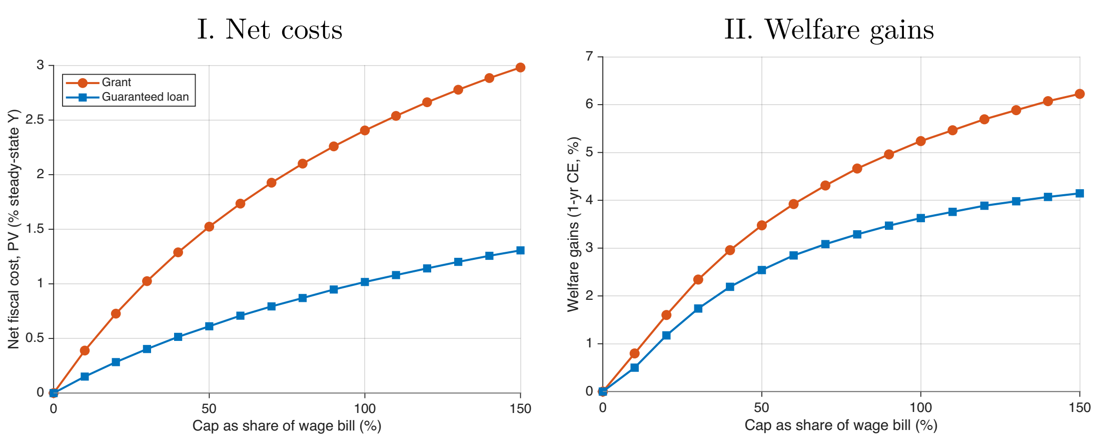

Notes: This figure presents the net costs and welfare gains of the grant and guaranteed-loan programs, varying the disbursement limit as a share of prior-year wage bill. The welfare gains are in 1-year consumption equivalent terms.

The right panel of Figure 9 presents the associated welfare gains. These gains are also concave and, over the caps considered, monotonically increasing. We note, however, that this is not generally true. At some point, the welfare gains start declining and eventually turn negative. This happens when a larger fraction of "saved" firms have an operating cost that exceeds their net present social value.

6 Conclusion

This paper studies how financial frictions distort the firm exit margin and quantifies the resulting aggregate costs. We develop a general equilibrium model of firm dynamics in which a single parameter governs the degree of financial frictions and the resulting excess exit. The model identifies a key informative statistic for the severity of financial frictions, the marginal propensity to exit with debt, which we estimate using confidential U.S. Census data and use to discipline the quantitative model. Financial frictions raise the exit rate from 5.0% to 9.3% and generate welfare losses of 3.6% in consumption-equivalent terms. These costs are amplified in financial crises, when collapsing borrowing capacity forces out viable but indebted firms, and are largely unchanged in productivity-driven recessions.

While the model speaks to the design of interventions that prevent excessive exit, a full normative evaluation is left for future research. We evaluate grants and guaranteed loans abstracting from the distortionary cost of the taxes that fund them and taking the targeting of support as given. Both are central to optimal policy, and targeting is likely first-order: the programs we study reach roughly a third of firms but prevent the exit of only one to two percent, so directing support toward viable but constrained firms, and away from firms that would exit efficiently, could deliver far larger gains per dollar.

A second direction for future research concerns the empirical object at the center of our analysis: the marginal propensity to exit with debt. Our model shows that this moment is central to calibrating the degree of financial frictions. We estimate it from confidential U.S. Census data, conditioning on a rich set of firm fundamentals. Given its importance, the moment invites further empirical work that measures it across settings and with sharper identification.

In summary, financial frictions impose substantial aggregate costs not only by distorting investment among continuing firms, the channel emphasized in prior work, but also, as we show, by inefficiently shaping which firms survive.

References

Albuquerque, Rui and Hugo A. Hopenhayn. 2004. "Optimal Lending Contracts and Firm Dynamics." *The Review of Economic Studies* 71 (2):285–315.

Arellano, Cristina, Yan Bai, and Patrick J Kehoe. 2019. "Financial Frictions and Fluctuations in Volatility." *Journal of Political Economy* 127 (5):2049–2103.

Ayres, Joao and Gajendran Raveendranathan. 2023. "Firm Entry and Exit during Recessions." *Review of Economic Dynamics* 47:47–66.

Buera, Francisco J, Joseph P Kaboski, and Yongseok Shin. 2011. "Finance and Development: A Tale of Two Sectors." *American Economic Review* 101 (5):1964–2002.

Castillo-Martinez, Laura. 2025. "Sudden Stops, Productivity, and the Exchange Rate." Tech. rep., CEPR Discussion Paper No. 20555.

Clementi, Gian Luca and Hugo A. Hopenhayn. 2006. "A Theory of Financing Constraints and Firm Dynamics." *The Quarterly Journal of Economics* 121 (1):229–265.

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

Crouzet, Nicolas and Neil R Mehrotra. 2020. "Small and Large Firms over the Business Cycle." *American Economic Review* 110 (11):3549–3601.

Crouzet, Nicolas and Fabrice Tourre. 2021. "Can the Cure Kill the Patient? Corporate Credit Interventions and Debt Overhang." Manuscript.

Cuciniello, Vincenzo, Claudio Michelacci, and Luigi Paciello. 2025. "Subsidizing Business Entry in Competitive Credit Markets." *Journal of Political Economy* 133 (11):3652–3711.

Daruich, Diego. 2026. "The Macroeconomic Consequences of Early Childhood Development Policies." *Journal of Political Economy* Forthcoming.

Dhingra, Swati and John Morrow. 2019. "Monopolistic Ccompetition and Optimum Product Diversity under Firm Heterogeneity." *Journal of Political Economy* 127 (1):196–232.

Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu. 2023. "How costly are markups?" *Journal of Political Economy* 131 (7):1619–1675.

Foster, Lucia, Cheryl Grim, and John Haltiwanger. 2016. "Reallocation in the Great Recession: Cleansing or Not?" *Journal of Labor Economics* 34 (1):293–331.

Gourinchas, Pierre-Olivier, Şebnem Kalemli-Özcan, Veronika Penciakova, and Nick Sander. 2025. "SME Failures under Large Liquidity Shocks: An Application to the COVID-19 Crisis." *Journal of the European Economic Association* 23 (2):431–480.

Hopenhayn, Hugo A. 1992. “Entry, Exit, and Firm Dynamics in Long Run Equilibrium.” *Econometrica* 60 (5):1127–1150.

Hsieh, Chang-Tai and Peter J Klenow. 2009. “Misallocation and Manufacturing TFP in China and India.” *The Quarterly journal of economics* 124 (4):1403–1448.

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

Karahan, Fatih, Benjamin Pugsley, and Ayşegül Şahin. 2024. “Demographic Origins of the Startup Deficit.” *American Economic Review* 114 (7):1986–2023.

Khan, Aubhik, Tatsuro Senga, and Julia K Thomas. 2021. “Default Risk and Aggregate Fluctuations in an Economy with Production Heterogeneity.” Manuscript.

Khan, Aubhik and Julia K Thomas. 2013. “Credit Shocks and Aggregate Fluctuations in an Economy with Production Heterogeneity.” *Journal of Political Economy* 121 (6):1055–1107.

Kochen, Federico. 2025. “Finance over the Life Cycle of Firms.” Manuscript.

Lee, Yoonsoo and Toshihiko Mukoyama. 2015. “Entry and Exit of Manufacturing Plants over the Business Cycle.” *European Economic Review* 77:20–27.

Leibovici, Fernando and David Wiczer. 2025. “Firm Exit and Liquidity: Evidence from the Great Recession.” Manuscript.

Lian, Chen and Yueran Ma. 2021. “Anatomy of Corporate Borrowing Constraints.” *The Quarterly Journal of Economics* 136 (1):229–291.

Manova, Kalina. 2013. “Credit Constraints, Heterogeneous Firms, and International Trade.” *Review of Economic Studies* 80 (2):711–744.

Midrigan, Virgiliu and Daniel Yi Xu. 2014. “Finance and Misallocation: Evidence from Plant-Level Data.” *American Economic Review* 104 (2):422–458.

Moll, Benjamin. 2014. “Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation?” *American Economic Review* 104 (10):3186–3221.

Moody’s Investors Service. 2026. “Annual Default Study: Corporate Default and Recovery Rates, 1920–2025.” Tech. rep., Moody’s Investors Service.

Ottonello, Pablo and Thomas Winberry. 2020. “Financial Heterogeneity and the Investment Channel of Monetary Policy.” *Econometrica* 88 (6):2473–2502.

———. 2024. “Capital, Ideas, and the Costs of Financial Frictions.” Working Paper.

Pugsley, Benjamin W., Petr Sedláček, and Vincent Sterk. 2021. “The Nature of Firm Growth.” *American Economic Review* 111 (2):547–579.

Quadrini, Vincenzo. 2009. “Entrepreneurship in Macroeconomics.” *Annals of Finance* 5 (3):295–311.

Restuccia, Diego and Richard Rogerson. 2008. “Policy Distortions and Aggregate Productivity with Heterogeneous Establishments.” *Review of Economic Dynamics* 11 (4):707–720.

Tauchen, George. 1986. “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions.” *Economics Letters* 20 (2):177–181.

A Model Appendix

A.1 Firm Profit Maximization Problem

Consider the static problem of an operating firm j that hires labor n_{jt} and chooses price p_{jt} to maximize profits:

$$\pi_{jt} = \max_{\{p_{jt}, y_{jt}\}} \; p_{jt} y_{jt} - \frac{y_{jt} w_t}{Z_t a_{jt}},$$
$$\text{s.t.} \quad y_{jt} = p_{jt}^{-\epsilon} Y_t.$$

The firm's optimal price is

$$p_{jt} = \frac{\epsilon}{\epsilon - 1} \frac{w_t}{Z_t a_{jt}}.$$

Households minimize the cost of purchasing the aggregate good Y_t, yielding the aggregate price index

$$P_t = \left(\int_0^{J_t} p_{jt}^{1-\epsilon} dj \right)^{\frac{1}{1-\epsilon}}.$$

Imposing the normalization $P_t = 1$ and substituting the optimal price gives

$$\frac{\epsilon}{\epsilon - 1} \frac{w_t}{Z_t \overline{A}_t} = 1 \quad \Rightarrow \quad w_t = \frac{\epsilon - 1}{\epsilon} Z_t \overline{A}_t. \tag{A.1}$$

Hence, the relative price of firm j is $p_{jt} = \overline{A}_t / a_{jt}$ and labor demand follows from the production function $y_{jt} = Z_t a_{jt} n_{jt}$:

$$n_{jt} = \left(\frac{a_{jt}}{\overline{A}_t} \right)^{\epsilon} \frac{1}{a_{jt}} \frac{Y_t}{Z_t}.$$

Aggregating over firms yields

$$Y_t = \overline{A}_t Z_t N_t, \tag{A.2}$$

where N_t is aggregate labor employed in production. These conditions jointly define the static decentralized equilibrium used below. The labor employed in direct production in each firm is then given by

$$n_{jt} = \left(\frac{a_{jt}}{\overline{A}_t} \right)^{\epsilon - 1} N_t. \tag{A.3}$$

The operating profits of firm j at time t are therefore given by

$$\pi(a_{jt}; \Omega_t) = \frac{1}{\epsilon} a_{jt}^{\epsilon - 1} \overline{A}_t^{2-\epsilon} N_t Z_t.$$

A.2 Social Planner's Problem and Efficient Allocation

This appendix characterizes the efficient allocation by solving the planner's problem in two steps: (i) the efficient allocation of labor across active firms, and (ii) the optimal exit and entry conditions.

A.2.1 Efficient Allocation Across Operating Firms

The planner allocates N_t units of labor across a measure J_t of firms with heterogeneous productivities $a_{jt}Z_t$ to maximize aggregate consumption,

$$\max_{\{c_{jt}\}} \left(\int_0^{J_t} c_{jt}^{\frac{\epsilon-1}{\epsilon}} dj \right)^{\frac{\epsilon}{\epsilon-1}} \quad \text{s.t.} \quad \int_0^{J_t} \frac{c_{jt}}{a_{jt}Z_t} dj = N_t.$$

The first-order condition, $c_{jt}^{-\frac{1}{\epsilon}} C_t^{\frac{1}{\epsilon}} = \lambda_t/(a_{jt}Z_t)$, implies $\frac{c_{jt}}{c_{kt}} = \left(\frac{a_{jt}}{a_{kt}}\right)^{\epsilon}$. Using the resource constraint and defining $\overline{A}_t \equiv \left(\int_0^{J_t} a_{jt}^{\epsilon-1} dj\right)^{1/(\epsilon-1)}$, we obtain

$$c_{jt} = a_{jt}^{\epsilon} \overline{A}_t^{1-\epsilon} Z_t N_t, \qquad C_t = \overline{A}_t Z_t N_t. \tag{A.4}$$

Dividing the resource constraint firm by firm gives the implied labor allocation,

$$n_{jt} = \frac{c_{jt}}{a_{jt}Z_t} = \left(\frac{a_{jt}}{\overline{A}_t}\right)^{\epsilon-1} N_t. \tag{A.5}$$

Hence, labor is allocated proportionally to relative productivity, and aggregate consumption equals aggregate productivity times effective labor.

A.2.2 Efficient Entry and Exit

The planner decides which firms operate and how many new firms enter, according to the net present value of their social contribution.

Flow social benefits and costs. Each firm affects aggregate output through two channels: its contribution to aggregate production and the labor it absorbs for fixed operations. The firm's marginal contribution to aggregate productivity is given by

$$\frac{1}{\epsilon-1} \overline{A}_t^{2-\epsilon} a_{jt}^{\epsilon-1},$$

and, multiplying by $Z_t N_t$, the corresponding contribution to aggregate output is

$$\frac{1}{\epsilon-1} \overline{A}_t^{2-\epsilon} Z_t N_t a_{jt}^{\epsilon-1}.$$

Conversely, operating the firm diverts labor $f_t/Z_t a_t$ from production, implying an output flow cost of

$$-\overline{A}_t \frac{f_t}{a_t}.$$

Recursive social value of the firm. Let $W_{FB}(a_t, f_t; \Omega_t)$ denote the firm's social value (in final consumption units):

$$W_{FB}(a_t, f_t; \Omega_t) = \frac{1}{\epsilon - 1}\overline{A}_t^{2-\epsilon} Z_t N_t a_t^{\epsilon-1} - \overline{A}_t \frac{f_t}{a_t} + \mathbb{E}_t\big[\mathcal{M}_{t,t+1}\widetilde{W}_{FB}(a_{t+1}, f_{t+1}; \Omega_{t+1})\big],$$

where $\widetilde{W}_{FB}(\cdot) \equiv \max\{W_{FB}(\cdot), 0\}$ is the pre-decision social value, and $\mathcal{M}_{t,t+1} = \beta \frac{C_t}{C_{t+1}}$ is the stochastic discount factor.

Efficient exit. The planner keeps a firm active if $W_{FB}(a_t, f_t; \Omega_t) \geq 0$, which implies

$$\frac{\overline{A}_t f_t}{a_t} \leq \frac{1}{\epsilon - 1}\overline{A}_t^{2-\epsilon} Z_t N_t a_t^{\epsilon-1} + \mathbb{E}_t\big[\mathcal{M}_{t,t+1}\widetilde{W}_{FB}(a_{t+1}, f_{t+1}; \Omega_{t+1})\big].$$

Hence, the efficient exit threshold is

$$\overline{f}_x^{FB}(a_t; \Omega_t) \equiv a_t Z_t \left[\frac{1}{\epsilon - 1}\overline{A}_t^{1-\epsilon} N_t a_t^{\epsilon-1} + \frac{1}{\overline{A}_t Z_t}\, \mathbb{E}_t\big[\mathcal{M}_{t,t+1}\widetilde{W}_{FB}(a_{t+1}, f_{t+1}; \Omega_{t+1})\big]\right]. \tag{A.6}$$

Efficient entry. The total entry cost is $\frac{1}{2} f_e M_{e,t}^2$ in effective units of labor. Since one effective unit corresponds to $1/Z_t$ units of raw labor, the marginal entrant diverts $f_e M_{e,t}/Z_t$ units of labor from production. Valued at the output per unit of labor, $\overline{A}_t Z_t$, the social cost of entry in output units is $f_e M_{e,t}\, \overline{A}_t$.

The planner expands entry until this equals the expected social value of a new firm:

$$f_e M_{e,t} \overline{A}_t = \int\int \widetilde{W}_{FB}(a_t, f_t; \Omega_t)\, dF(a_t)\, dG(f_t). \tag{A.7}$$

In sum, the efficient labor allocation across firms (A.5), the exit threshold (A.6), and the free-entry condition (A.7) jointly characterize the efficient allocation.

A.3 Proofs

A.3.1 Proof of Proposition 1

To prove that the decentralized and efficient allocations coincide, we will show that the equilibrium conditions are equivalent. In particular, we show labor is distributed efficiently in the market allocation according to (A.5), the exit threshold is equivalent to (A.6), and free entry satisfies (A.7).

No misallocation in production labor We start by noting that labor is efficiently allocated in the market allocation. That is, the equation that pins down the allocation of production labor across firms in the market allocation (A.3) is identical to the efficient one (A.5).

Exit thresholds With $\beta = 0$, the stochastic discount factor $\mathcal{M}_{t,t+1} = 0$, so lenders receive no expected return from extending credit. Consequently, the debt price satisfies $q(a_t, b_t; \Omega_t) = 0$ for all firms, implying a zero borrowing limit and no outstanding debt in equilibrium. With $b_{t-1} = 0$ and no continuation value,

$$\widetilde{V}(a_t, \frac{w_t f_t}{Z_t a_t}; \Omega_t) = \max\left\{\pi(a_t;\Omega_t) - \frac{w_t f_t}{Z_t a_t}, 0\right\}, \qquad \widetilde{W}_{FB}(a_t, f_t; \Omega_t) = \max\left\{\tfrac{1}{\epsilon-1}\overline{A}_t^{2-\epsilon} Z_t N_t a_t^{\epsilon-1} - \overline{A}_t \tfrac{f_t}{a_t}, 0\right\}.$$

Substituting (2.1) and (A.1) into the first expression yields the pointwise proportionality

$$\widetilde{V}(a_t, \frac{w_t f_t}{Z_t a_t}; \Omega_t) = \frac{\epsilon - 1}{\epsilon} \widetilde{W}_{FB}(a_t, f_t; \Omega_t),$$

so the zero-value cutoff is the same in both problems. Hence the decentralized and planner exit thresholds *coincide firm by firm* for any given Ω_t.

Free entry From the main text and Appendix A.2, the decentralized and planner free-entry conditions are

$$\text{CE: } \frac{f_e M_{e,t} w_t}{Z_t} = \iint \widetilde{V}(a, \frac{w_t f}{Z_t a}; \Omega_t)\, dF(a)\, dG(f), \qquad \text{SP: } f_e M_{e,t}\, \overline{A}_t = \iint \widetilde{W}_{FB}(a, f; \Omega_t)\, dF(a)\, dG(f).$$

Using $\widetilde{V} = \frac{\epsilon-1}{\epsilon} \widetilde{W}_{FB}$ and $w_t = \frac{\epsilon-1}{\epsilon} \overline{A}_t Z_t$, the CE condition transforms into the SP condition exactly. Thus, the mass of entrants $M_{e,t}$ and the implied aggregate state Ω_t (hence $\overline{A}_t, N_t$) are the same in both economies.

Since exit rules coincide firm by firm, free entry pins down the same mass of entrants, and production labor is efficiently allocated, the decentralized equilibrium is identical to the social planner's one. □

A.3.2 Proof of Proposition 2

To prove that the decentralized and efficient allocations coincide, we will show that the equilibrium conditions are equivalent. In particular, we show labor is distributed efficiently in the market allocation according to (A.5), the exit threshold is equivalent to (A.6), and free entry satisfies (A.7).

No misallocation in production labor We start by noting that labor is efficiently allocated in the market allocation. That is, the equation that pins down the allocation of production labor across firms in the market allocation (A.3) is identical to the efficient one (A.5).

Exit thresholds. Plugging the expression for profits (2.1) and the real wage (A.1), we have that the decentralized equilibrium exit threshold under $\kappa = 0$ is

$$f_x(a_t, b_{t-1}; \Omega_t) = \bar{f}(a_t, 0; \Omega_t) = a_t Z_t \left[\frac{1}{\epsilon - 1} a_t^{\epsilon-1} \overline{A}_t^{1-\epsilon} N_t + \frac{\epsilon}{\epsilon - 1} \frac{1}{Z_t \overline{A}_t} \bar{L}(a_t; \Omega_t)\right].$$

To prove that the threshold rule coincides with that of the planner, we need to show that

$$\bar{L}(a_t; \Omega_t) = \frac{\epsilon - 1}{\epsilon} \mathbb{E}_t \left[\mathcal{M}_{t,t+1} \tilde{W}_{FB}\left(a_{t+1}, f_{t+1}; \Omega_{t+1}\right)\right].$$

We proceed in three steps. First, we show that the borrowing amount is weakly increasing in b_t. Second, we show that the borrowing limit equals the expected private continuation value of the firm. Finally, we show that the expected private and social continuation values coincide, implying identical exit thresholds in the decentralized and planner economies.

(i) Borrowing amount is increasing in b Recall that the borrowing limit is given by

$$\overline{L}(a_t;\Omega_t) = \max_b q(a_t,b;\Omega_t)b,$$

where

$$\begin{aligned} q(a_t,b_t;\Omega_t) =& \mathbb{E}_t\left[\mathcal{M}_{t,t+1}G\left(\bar{f}\left(a_{t+1},b_t;\Omega_{t+1}\right)\right)\right] \\ &+ \mathbb{E}_t\left[\mathcal{M}_{t,t+1}\frac{\widetilde{V}\left((1-\kappa)a_{t+1},\frac{f_{t+1}w_{t+1}}{Z_{t+1}(1-\kappa)a_{t+1}};\Omega_{t+1}\right)}{b_t}\mathbb{1}\left[f_{t+1}>\bar{f}\left(a_{t+1},b_t;\Omega_{t+1}\right)\right]\right]. \end{aligned}$$

Define the borrowing amount when debt issuance is equal to b_t to be $L(a_t,b_t;\Omega_t)$. So that $\overline{L}(a_t;\Omega_t) = \max_b L(a_t,b;\Omega_t)$. Using the equation above, when $\kappa=0$ we have

$$\begin{aligned} L(a_t,b_t;\Omega_t) =& \mathbb{E}_t\left[\mathcal{M}_{t,t+1}G\left(\bar{f}\left(a_{t+1},b_t;\Omega_{t+1}\right)\right)b_t\right] \\ &+ \mathbb{E}_t\left[\mathcal{M}_{t,t+1}\widetilde{V}\left(a_{t+1},\frac{f_{t+1}w_{t+1}}{Z_{t+1}a_{t+1}};\Omega_{t+1}\right)\mathbb{1}\left[f_{t+1}>\bar{f}\left(a_{t+1},b_t;\Omega_{t+1}\right)\right]\right]. \end{aligned}$$

Differentiating with respect to b_t we obtain,

$$\begin{aligned} \frac{\partial L(a_t,b_t;\Omega_t)}{\partial b_t} =& \mathbb{E}_t\left[\mathcal{M}_{t,t+1}G\left(\bar{f}\left(a_{t+1},b_t;\Omega_{t+1}\right)\right)\right] \\ &+ \mathbb{E}_t\left[\mathcal{M}_{t,t+1}g\left(\bar{f}\left(a_{t+1},b_t;\Omega_{t+1}\right)\right)\frac{Z_{t+1}a_{t+1}}{w_{t+1}}\left[\widetilde{V}\left(a_{t+1},\frac{\bar{f}\left(a_{t+1},b_t;\Omega_{t+1}\right)w_{t+1}}{Z_{t+1}a_{t+1}};\Omega_{t+1}\right)-b_t\right]\right], \end{aligned}$$

where we've used $\frac{\partial\bar{f}(a_{t+1},b_t;\Omega_{t+1})}{\partial b_t} = -\frac{Z_{t+1}a_{t+1}}{w_{t+1}}$. From the equation above, we can see that if

$$\widetilde{V}\left(a_{t+1},\frac{\bar{f}\left(a_{t+1},b_t;\Omega_{t+1}\right)w_{t+1}}{Z_{t+1}a_{t+1}};\Omega_{t+1}\right) \geq b_t, \tag{A.8}$$

for all a_{t+1} and Ω_{t+1}, then $\frac{\partial L(a_t,b_t;\Omega_t)}{\partial b_t}$ is positive. By definition, we have

$$\widetilde{V}\left(a_{t+1},b_t+\frac{\bar{f}\left(a_{t+1},b_t;\Omega_{t+1}\right)w_{t+1}}{Z_{t+1}a_{t+1}};\Omega_{t+1}\right) \geq 0, \tag{A.9}$$

because the firm chooses not to default. Note that since the firm chooses not to default $\tilde{V}(\cdot) = V(\cdot)$

which is defined in (2.4):

$$V\left(a_t, x_t; \Omega_t\right) = \max_{\{b_t, d_t\}} d_t + \mathbb{E}_t\left[\mathcal{M}_{t,t+1}\widetilde{V}\left(a_{t+1}, b_t + \frac{f_{t+1}w_{t+1}}{Z_{t+1}a_{t+1}}; \Omega_{t+1}\right)\right],$$
$$\text{s.t.} \quad d_t = \pi(a_t; \Omega_t) - x_t + q(a_t, b_t; \Omega_t)b_t$$
$$d_t \geq 0$$

Using the envelope theorem, we have that

$$\frac{\partial V\left(a_t, x_t; \Omega_t\right)}{\partial x_t} = -\lambda,$$

where λ is the Lagrange multiplier on the dividend constraint. The FOC with respect to d_t implies that

$$\lambda = 1 + \nu,$$

where $\nu \geq 0$ is the Lagrange multiplier on the non-negativity of dividends constraint. So we obtain $\frac{\partial V(a_t,x_t;\Omega_t)}{\partial x_t} \leq -1$ for all $\{a_t, x_t, \Omega_t\}$. We can now use this derivative to obtain

$$\begin{aligned}
&\widetilde{V}\left(a_{t+1}, \frac{\bar{f}\left(a_{t+1}, b_t; \Omega_{t+1}\right) w_{t+1}}{Z_{t+1}a_{t+1}}; \Omega_{t+1}\right) \\
&= V\left(a_{t+1}, \frac{\bar{f}\left(a_{t+1}, b_t; \Omega_{t+1}\right) w_{t+1}}{Z_{t+1}a_{t+1}}; \Omega_{t+1}\right) \\
&= V\left(a_{t+1}, b_t + \frac{\bar{f}\left(a_{t+1}, b_t; \Omega_{t+1}\right) w_{t+1}}{Z_{t+1}a_{t+1}}; \Omega_{t+1}\right) - \int_0^{b_t} \frac{\partial V\left(a_{t+1}, b + \frac{\bar{f}(a_{t+1}, b_t; \Omega_{t+1}) w_{t+1}}{Z_{t+1}a_{t+1}}; \Omega_{t+1}\right)}{\partial x_t} db \\
&\geq V\left(a_{t+1}, b_t + \frac{\bar{f}\left(a_{t+1}, b_t; \Omega_{t+1}\right) w_{t+1}}{Z_{t+1}a_{t+1}}; \Omega_{t+1}\right) + b_t \\
&\geq b_t
\end{aligned}$$

where the second-to-last inequality follows from $\frac{\partial V(a_t,x_t;\Omega_t)}{\partial x_t} \leq -1$ and the last inequality follows from (A.9). Thus, condition (A.8) is satisfied and $\frac{\partial L(a_t,b_t;\Omega_t)}{\partial b_t} \geq 0$.

(ii) Borrowing limit equals private continuation value As a result, we have that $\overline{L}(a_t; \Omega_t) = \lim_{b_t \to \infty} L(a_t, b_t; \Omega_t)$. Plugging into the $L(\cdot)$ function, we obtain

$$\overline{L}(a_t; \Omega_t) = \mathbb{E}_t\left[\mathcal{M}_{t,t+1}\widetilde{V}\left(a_{t+1}, \frac{f_{t+1}w_{t+1}}{Z_{t+1}a_{t+1}}; \Omega_{t+1}\right)\right] \tag{A.10}$$

because $\bar{f}\left(a_{t+1}, b_t; \Omega_{t+1}\right) < 0$ when $b_t \to \infty$. Note that when $\kappa = 0$, the firm exits if and only if $f > \bar{f}\left(a_{t+1}, 0; \Omega_{t+1}\right)$. This follows from the fact that the firm's productivity remains unchanged after default and that $\bar{f}\left(a_{t+1}, b_t; \Omega_{t+1}\right)$ decreases with b_t. So if the firm owner would choose to exit without any debt, the firm would also choose to default with any level of positive debt.

(iii) Private and social continuation values coincide The final step to verify that the decentralized and efficient exit thresholds coincide is to show that the expected continuation value of the firm is identical in both problems. Specifically, we need to establish that

$$\mathbb{E}_t\left[\frac{\widetilde{V}\left(a_{t+1}, \frac{f_{t+1}w_{t+1}}{Z_{t+1}a_{t+1}}; \Omega_{t+1}\right)}{\overline{A}_{t+1}Z_{t+1}N_{t+1}}\right] = \frac{\epsilon-1}{\epsilon}\mathbb{E}_t\left[\frac{\widetilde{W}_{FB}\left(a_{t+1}, f_{t+1}; \Omega_{t+1}\right)}{\overline{A}_{t+1}Z_{t+1}N_{t+1}}\right].$$

where we have substituted for the stochastic discount factor $\mathcal{M}_{t,t+1}$ and conditioned on Ω_t. We will show that

$$\frac{\epsilon}{\epsilon-1}\frac{\widetilde{V}\left(a_{t+1}, \frac{f_{t+1}w_{t+1}}{Z_{t+1}a_{t+1}}; \Omega_{t+1}\right)}{\overline{A}_{t+1}Z_{t+1}N_{t+1}} = \frac{\widetilde{W}_{FB}\left(a_{t+1}, f_{t+1}; \Omega_{t+1}\right)}{\overline{A}_{t+1}Z_{t+1}N_{t+1}}. \tag{A.11}$$

Under standard regularity conditions—$\beta < 1$ and a bounded productivity process, which imply that V and $\widetilde{V}$ are bounded—the Bellman operator is a contraction with a unique fixed point. The argument below uses a guess-and-verify approach: we conjecture that $b_t \to \infty$ is optimal when $\kappa = 0$, derive the resulting value function, and then verify that the derived function satisfies the original Bellman equation. In this limiting case, the firm effectively transfers ownership to creditors each period; creditors accept this because $\kappa = 0$ implies no productivity loss following default, so they recover the full continuation value of the firm. The monotonicity of $L(a, b)$ in b established in step (i), combined with boundedness of $\widetilde{V}$, ensures that $\lim_{b\to\infty} L(a, b)$ exists and is finite. In that case, we have that

$$V\left(a_t, x_t; \Omega_t\right) = \pi(a_t; \Omega_t) - x_t + \mathbb{E}_t\left[\mathcal{M}_{t,t+1}\widetilde{V}\left(a_{t+1}, \frac{f_{t+1}w_{t+1}}{Z_{t+1}a_{t+1}}; \Omega_{t+1}\right)\right]. \tag{A.12}$$

where we have already used the expression for the borrowing limit (A.10).

Let's define $V^v(a_t, f_t; \Omega_t) \equiv \frac{\epsilon}{\epsilon-1}\frac{V\left(a_t, \frac{f_t w_t}{Z_t a_t}; \Omega_t\right)}{\overline{A}_t Z_t N_t}$ and $W^w_{FB}\left(a_t, f_t; \Omega_t\right) = \frac{W_{FB}(a_t, f_t; \Omega_t)}{\overline{A}_t Z_t N_t}$. Defining $V^v(\cdot)$ and $W^w(\cdot)$ in this way, it is sufficient to show that they coincide to verify condition (A.11).

Using equation (A.12) and the definition of $V^v(a_t, f_t; \Omega_t)$

$$V^v\left(a_t, f_t; \Omega_t\right) = \frac{\epsilon}{\epsilon-1}\frac{1}{Z_t\overline{A}_tN_t}\left(\pi(a_t; \Omega_t) - \frac{f_tw_t}{Z_ta_t}\right) + \beta\mathbb{E}_t\left[\frac{\epsilon}{\epsilon-1}\frac{1}{Z_{t+1}\overline{A}_{t+1}N_{t+1}}\widetilde{V}\left(a_{t+1}, \frac{f_{t+1}w_{t+1}}{Z_{t+1}a_{t+1}}; \Omega_{t+1}\right)\right].$$

Given $\widetilde{V}^v\left(a_t, f_t; \Omega_t\right) = \max\{V^v\left(a_t, f_t; \Omega_t\right), 0\}$ and the expressions for profits (2.1) and the real wage (A.1), this simplifies to

$$V^v\left(a_t, f_t; \Omega_t\right) = \frac{1}{Z_t\overline{A}_tN_t}\left(\frac{1}{\epsilon-1}a_t^{\epsilon-1}\overline{A}_t^{2-\epsilon}N_tZ_t - \frac{f_t\overline{A}_t}{a_t}\right) + \beta\mathbb{E}_t\left[\widetilde{V}^v\left(a_{t+1}, f_{t+1}; \Omega_{t+1}\right)\right], \tag{A.13}$$

Using equation (3.3) and the definition of $W^w_{FB}\left(a_t, f_t; \Omega_t\right)$, it follows that

$$W^w_{FB}\left(a_t, f_t; \Omega_t\right) = \frac{1}{Z_t\overline{A}_tN_t}\left(\frac{1}{\epsilon-1}a_t^{\epsilon-1}\overline{A}_t^{2-\epsilon}N_tZ_t - \frac{f_t\overline{A}_t}{a_t}\right) + \beta\mathbb{E}_t\left[\widetilde{W}^w_{FB}\left(a_{t+1}, f_{t+1}; \Omega_{t+1}\right)\right] \tag{A.14}$$

where $\widetilde{W}^w_{FB}(a_t, f_t; \Omega_t) = \max\{W^w_{FB}(a_t, f_t; \Omega_t), 0\}$.

Note that equations (A.13)–(A.14) are equivalent, so condition (A.11) holds. This confirms our conjecture that taking $b_t \to \infty$ indeed maximizes the firm's private value, as the owner obtains the complete-markets value of the firm. Since the borrowing limit equals the firm's private continuation value, which in turn equals its social continuation value, the decentralized and planner exit thresholds coincide.

Free entry. From the main text and Appendix A.2, the decentralized and planner free-entry conditions are

$$\text{CE: } \frac{f_e M_{e,t} w_t}{Z_t} = \iint \widetilde{V}\left(a, \frac{f w_t}{Z_t a}; \Omega_t\right) dF(a)\, dG(f), \qquad \text{SP: } f_e M_{e,t}\, \overline{A}_t = \iint \widetilde{W}_{FB}(a, f; \Omega_t)\, dF(a)\, dG(f).$$

Using $w_t = \frac{\epsilon-1}{\epsilon}\overline{A}_t Z_t$ and equation (A.11), the CE condition transforms exactly into the SP condition. Hence $M_{e,t}$ and the implied aggregate state Ω_t coincide across the two economies.

Since exit rules coincide firm by firm and free entry pins down the same mass of entrants, the survivor operator and aggregates $(\overline{A}_t, N_t)$ coincide. Together with identical aggregate accounting, the decentralized allocation equals the efficient allocation when $\kappa = 0$. □

A.3.3 Proof of Lemma 1

From equation (3.7),

$$P_x(b) = 1 - G\Big(\max\{\pi + \bar{L} - b,\, (1-\kappa)\pi + \bar{L}^\kappa\}\Big).$$

Let

$$A \equiv \pi + \bar{L} \quad \text{and} \quad B \equiv (1-\kappa)\pi + \bar{L}^\kappa,$$

so that the cutoff debt level is $\bar{b} = A - B$.

When $b < \bar{b}$, we have $A - b > B$, and therefore

$$P_x(b) = 1 - G(A - b).$$

Differentiating with respect to b yields

$$P_x'(b) = -G'(A-b)\,(-1) = g(A-b) = g(\pi + \bar{L} - b).$$

When $b > \bar{b}$, we have $A - b < B$, so $P_x(b) = 1 - G(B)$ is constant in b, implying $P_x'(b) = 0$. At $b = \bar{b}$, the left-hand derivative equals $g(B)$ while the right-hand derivative is zero, generating a kink at the cutoff. □

A.3.4 Proof of Proposition 3

The proof proceeds in three steps. First, we show that the borrowing limit $\bar{L}$ is weakly decreasing in the degree of financial frictions, κ. Second, we establish that the cutoff debt level $\bar{b}$ is increasing in κ.

Finally, we combine both results to show that the marginal propensity to exit with debt, $P_x'(b)$, rises with κ for any given debt level.

(i) $\bar{L}$ is decreasing in κ. For notational convenience, set

$$A(\pi_0,\kappa) \equiv (1-\kappa)\pi_0, \qquad z \equiv \pi_0 - b \in [0,\pi_0].$$

Multiplying (3.5) by $b = \pi_0 - z$ and changing variables yields the equivalent program

$$\bar{L}(\pi_0,\kappa) = \beta \max_{z\in[0,\pi_0]} \Phi(z;\pi_0,\kappa), \quad \text{where} \quad \Phi(z;\pi_0,\kappa) = (\pi_0 - z)\, G(z) \ + \int_z^{A(\pi_0,\kappa)} (A(\pi_0,\kappa) - f)\, dG(f),$$

with the convention that the integral is 0 when $z \geq A(\pi_0,\kappa)$. An integration-by-parts computation (using $dG = g\,df$ on $[\underline{f},\infty)$) gives, for $z \leq A(\pi_0,\kappa)$,

$$\Phi(z;\pi_0,\kappa) = \kappa\,\pi_0\, G(z) + \int_z^{A(\pi_0,\kappa)} G(f)\, df.$$

Consequently, the envelope (one-sided) derivatives of $\bar{L}$ are

$$\frac{\partial \bar{L}}{\partial \kappa}(\pi_0,\kappa) = \begin{cases} \beta\,\pi_0\big[G(z^*) - G(A(\pi_0,\kappa))\big] \ \leq 0, & \text{if } z^* \leq A(\pi_0,\kappa), \\ 0, & \text{if } z^* \geq A(\pi_0,\kappa), \end{cases} \tag{A.15}$$

where $z^* = z^*(\pi_0,\kappa)$ is any maximizer of $\Phi(\cdot;\pi_0,\kappa)$. Hence $\bar{L}$ is (weakly) decreasing in κ.

(ii) $\bar{b}$ is increasing in κ. Next, let's turn to the derivative of $\bar{L}$ with respect to profits:

$$\frac{\partial \bar{L}}{\partial \pi_0}(\pi_0,\kappa) = \begin{cases} \beta\big[\kappa\, G(z^*) + (1-\kappa)\, G(A(\pi_0,\kappa))\big] \ \geq 0, & \text{if } z^* \leq A(\pi_0,\kappa), \\ \beta\, G(z^*) \ \geq 0, & \text{if } z^* \geq A(\pi_0,\kappa). \end{cases} \tag{A.16}$$

Define

$$\bar{b}(\kappa) = \kappa\pi + \bar{L}(\pi,\kappa) - \underbrace{\bar{L}\big((1-\kappa)\pi,\kappa\big)}_{\equiv \bar{L}^{\kappa}}.$$

Differentiating gives

$$\bar{b}'(\kappa) = \pi + \frac{\partial \bar{L}}{\partial \kappa}(\pi,\kappa) - \Big[\frac{\partial \bar{L}}{\partial \kappa}\big((1-\kappa)\pi,\kappa\big) - \pi\,\frac{\partial \bar{L}}{\partial \pi_0}\big((1-\kappa)\pi,\kappa\big)\Big]. \tag{A.17}$$

From (A.15) with $\pi_0 = \pi$ and $A(\pi,\kappa) = (1-\kappa)\pi$,

$$\frac{\partial \bar{L}}{\partial \kappa}(\pi,\kappa) \ \geq \ -\beta\,\pi\, G\big((1-\kappa)\pi\big).$$

Also, by (A.15) and (A.16) at $\pi_0 = (1-\kappa)\pi$,

$$-\frac{\partial \bar{L}}{\partial \kappa}((1-\kappa)\pi, \kappa) \;\geq\; 0, \qquad \frac{\partial \bar{L}}{\partial \pi_0}((1-\kappa)\pi, \kappa) \;\geq\; 0.$$

Plugging these bounds into (A.17) and dropping the non-negative terms yields the uniform lower bound

$$\bar{b}'(\kappa) \;\geq\; \pi\Big[1 - \beta\, G\big((1-\kappa)\pi\big)\Big] \;>\; 0, \tag{A.18}$$

so $\bar{b}$ is strictly increasing in κ.

(iii) Putting the pieces together. By Lemma 1,

$$P_x'(b;\kappa) = g\big(\pi + \bar{L}(\pi,\kappa) - b\big)\, \mathbf{1}\{\, b < \bar{b}(\kappa)\,\}.$$

By (i), $\bar{L}(\pi,\kappa)$ decreases with κ, while by (ii), $\bar{b}(\kappa)$ increases with κ. Thus, higher κ both expands the region where $P_x'(b;\kappa) > 0$ and lowers the argument $\pi + \bar{L}(\pi,\kappa) - b$. Since g is decreasing, this implies that both the support and the magnitude of $P_x'(b;\kappa)$ increase with κ. □

B Data Appendix: Italian Firm-Level Data

This appendix describes the construction of the Italian firm-level dataset used in Section 4.3. The data come from the Orbis database, produced by Bureau van Dijk (Moody's Analytics), the predominant source for multi-country firm-level analysis. Its coverage is particularly strong for European countries, where mandatory filing requirements ensure broad representation of small and medium-sized firms alongside larger enterprises. Orbis reports annual balance sheet and profit-and-loss data, allowing us to construct measures of firm financial health that mirror those available in the QFR.

We focus on Italian manufacturing firms (NACE Rev. 2, divisions 10–33) over the period 2002–2016. The choice of country is guided by data quality: Italy has high coverage rates relative to national aggregates. We exclude observations with negative total assets, negative sales, or internally inconsistent balance sheet identities. Financial variables are deflated using the Italian GDP deflator (base year 2010 = 100) from Eurostat. Our main measure of firm indebtedness, the net short-term debt-to-sales ratio, is defined identically to the US analysis. All continuous financial variables are winsorized at the 5th and 95th percentiles within each year. A firm is classified as exiting in year t if its Orbis status indicator records a terminal event—bankruptcy, dissolution, or liquidation. The exit is dated to the year following the firm's last year of reported financial statements.

To assess how sample composition affects the estimated debt–exit relationship, we construct a QFR-comparable subsample: firms whose total assets exceed €311,000 in constant 2010 euros, the Italian equivalent of the QFR's $250,000 (2000 USD) asset cutoff. A firm enters this subsample from the first year in which it crosses the threshold and remains thereafter. After cleaning, the sample comprises 229,637 unique firms and approximately 1.97 million firm–year observations.

We validate the representativeness of our sample along two dimensions. Table A.3 reports the ratio of Orbis sample aggregates to EU KLEMS national accounts for Italian manufacturing. Coverage ranges between 49% and 82% for employment, 72% and 91% for the wage bill, and 79% and 95% for output, with coverage generally improving over time. Table A.4 compares the distribution of economic activity across firm-size classes with Eurostat's Structural Business Statistics for the year 2012. The sample closely matches the actual distribution, particularly for medium and large firms; small firms (fewer than 10 employees) are slightly underrepresented, as expected. Table A.5 reports summary statistics on firm fundamentals and financial characteristics.

C Computational Appendix

This appendix describes the numerical algorithm we use to solve the model. We first describe the discretization of the state space, then the algorithm for the stationary equilibrium, the calibration procedure, the algorithm for perfect-foresight transition dynamics, and finally the implementation of the policy experiments. Throughout, it is convenient to collect the aggregate variables that scale a firm's flow profits in equation (2.1) into a single demand index D_t, so that profits are proportional to $a^{\epsilon-1}D_t$ and the real wage follows from the household's optimality conditions. In the stationary equilibrium, solving for the aggregate allocation is equivalent to solving for the scalar D.

C.1 State Space and Discretization

The idiosyncratic state of a firm consists of its permanent productivity a^p, its transitory productivity a^{tr}, and its beginning-of-period debt obligations b. We discretize this state space as follows.

Productivity. The permanent component takes $N_{a^p} = 3$ values, equally spaced in logs at $\{-2.5\sigma_a^p, 0, 2.5\sigma_a^p\}$ around the (de-meaned) mean, with entry probabilities given by the corresponding intervals of the normal distribution. The transitory component is discretized on a grid of $N_{a^{tr}} = 71$ points using the method of Tauchen (1986), with the grid spanning three unconditional standard deviations of the AR(1) process. The same procedure delivers the transition matrix over transitory productivity and the distribution of entrants' initial draws.

Debt. Debt obligations lie on a grid of $N_b = 150$ points on $[0, \bar{b}]$. The grid is polynomial of degree three, $b_i = (i/N_b)^3\,\bar{b}$, so that points are concentrated near zero where most of the mass of firms lies, and includes an exact zero (the state of entrants and of firms taken over after default).

Operating costs. We do not discretize the fixed operating cost f. Because the repayment, takeover, and exit decisions all take a threshold form in f (equations (2.3) and (2.6)), all objects that require integrating over f—repayment and default probabilities, expected operating-cost payments conditional on surviving, and the distribution of next-period debt obligations—have closed forms under the Pareto distribution G, which we use directly. In particular, for each state and each borrowing choice, the probability of landing on each point of the next-period obligation grid is computed as a difference of the Pareto CDF evaluated at the implied f-thresholds.

Default. The productivity loss upon default maps $\log a^{tr}$ into $\log a^{tr} + \log(1-\kappa)$, which generally falls between grid points. We represent this mapping as a Markov matrix that splits each grid point's mass between the two adjacent grid points with linear-interpolation weights. Objects evaluated at the degraded productivity level (the value of the firm under new ownership and its borrowing terms, which enter the debt-pricing equation (2.5)) are computed by spline interpolation of the corresponding functions along the transitory-productivity dimension.

C.2 Stationary Equilibrium

For a given demand index D, the firm's problem and the debt-pricing schedule form a joint fixed point: the schedule $q(a, b')$ depends on next period's repayment thresholds and on the value of defaulted firms, while values and thresholds depend on the schedule. We solve this fixed point by value function iteration:

0) Start with a guess for the value function on the grid, V. We use the present value of flow profits, $V = \pi(a)/(1-\beta)$.

1) Given the implied continuation values, update the debt-pricing schedule $q(a, b')$ from the intermediaries' break-even condition (2.5): the expected repayment probability follows from the Pareto thresholds, and the recovery value is the expected value of the firm under new ownership, evaluated by interpolation at $(1-\kappa)a^{tr}$. To enhance stability, we dampen the update of the schedule, with a weight of 0.3 on the new schedule. The borrowing limit is $\bar{L}(a) = \max_{b'} q(a, b')b'$ on the grid.

2) Given the schedule, compute the repayment threshold $\bar{f}(a, b)$ from equation (2.3) and the exit threshold from equation (2.6), and update the value function: current dividends use the closed-form conditional expectation of f given survival, and continuation values integrate over the transition matrices for productivity and obligations.

3) If $\max |V^{new} - V|$ is below tolerance (10^{-10} in the baseline), stop; otherwise dampen ($V = 0.3V^{new} + 0.7V$) and return to step 1.

With the converged policies, we compute the stationary distribution $\Lambda(a, b)$ by iterating on the law of motion (2.7), normalizing the measure of entrants to one, until the maximal change is below 10^{-12}. The repay and default-and-continue branches use their respective transition matrices; entrants are injected at zero debt with their initial productivity distribution.

Two equilibrium conditions remain: free entry and labor-market clearing. In the calibration we exploit the model's homogeneity: given D, the free-entry condition and the labor-market-clearing condition jointly pin down the measure of entrants M_e and the entry-cost parameter f_e in closed form, given the value of entry V_e and total labor demand per entrant implied by the normalized stationary distribution. The calibrated D is the value at which the implied moments match the data (next subsection), and we verify that markets clear at the calibrated parameter vector. For counterfactual steady states—most importantly the frictionless economy with $\kappa = 0$—we hold f_e fixed at its calibrated value and solve for the new equilibrium D by bisection on the labor-market-clearing residual, with tolerance 10^{-8}; each evaluation of the residual requires re-solving steps 0–3 and the stationary distribution at the candidate D.

C.3 Calibration

We calibrate the nine internal parameters by simulated method of moments. Because each evaluation of the model is costly, we proceed in three steps. First, we evaluate the model on quasi-random Sobol draws over a hypercube of the parameter space, computing at each draw the full set of targeted moments (exit rates by age, computed by pushing a cohort of entrants through the stationary transition operator; average startup size; employment growth of surviving three-year-old firms by initial size; and the exit-on-debt regression coefficient, computed by weighted least squares on model-generated state-level data with productivity fixed effects, weighting each state by its mass in the stationary distribution). We iteratively narrow the hypercube around the best-fitting region; the final round uses 1,000 draws. Second, we fit a third-order polynomial response surface mapping parameters to moments on the final round of draws, weighting draws by the inverse of their distance to the data moments. Third, we minimize the moment distance over the fitted surface, which delivers the calibrated parameter vector. We verify that the model solved at the calibrated vector reproduces the targeted moments.

C.4 Transition Dynamics

All experiments in Section 5 are unexpected ("MIT") shocks followed by perfect-foresight transitions back to a stationary equilibrium. A transition is a path $\{D_t\}_{t=1}^{T}$ such that, given the path, firm values and policies are consistent with backward induction from the terminal steady state, the distribution of firms evolves consistently from the initial distribution, and labor markets clear in every period. The horizon T is chosen large enough that the economy has converged ($T = 100$ for the transitory shocks, $T = 150$–200 for the policy experiments and the removal of financial frictions); we verify that results are insensitive to T.

The algorithm iterates on the path of the demand index:

0) Guess a path $\{D_t\}$. For shocks that revert to the initial steady state we initialize at the steady-state value; for the transition to the frictionless economy we interpolate linearly between the initial and terminal values.

1) *Backward pass.* Starting from the terminal steady-state value function, iterate the firm's problem backward from $t = T - 1$ to $t = 1$. Each step updates the period-t debt-pricing schedule, repayment probabilities, recovery values, and value function, given period-$t+1$ objects, exactly as in step 1–2 of the stationary algorithm.

2) *Forward pass.* Starting from the initial stationary distribution, iterate the distribution forward. In each period, the value of entry determines the measure of entrants from the free-entry condition, the transition matrices implied by the period-t policies update the distribution, and labor-market clearing delivers the implied demand index $\widetilde{D}_t$ along with aggregate consumption, output, entry, exit, and default rates.

3) If $\max_t |\widetilde{D}_t - D_t|/D_t$ is below tolerance (10^{-3}; 10^{-4} for the policy-size sweep in Figure 9), stop.

Otherwise update D_t toward $\widetilde{D}_t$ with a damping weight that increases linearly across iterations (from 0.15 to 0.35 for the shock experiments) and return to step 1.

Aggregate shocks enter the backward and forward passes directly: aggregate productivity Z_t scales the effective demand index and converts operating and entry costs into effective labor units, while a time-varying κ_t changes the default operator and the recovery values period by period.

C.5 Policy Experiments

The grant and guaranteed-loan programs of Section 5 require two extensions. First, both programs cap transfers by the firm's *prior-year* wage bill and restrict eligibility by prior-year employment, which are not part of the recursive state. We therefore construct, once, an augmented distribution of firms at the policy date that records both the current state and the productivity in the preceding period, from which prior-year wage bills and eligibility shares are computed. Second, the transfer received by a firm depends on its cash-flow shortfall, which depends on the same operating-cost draw f that determines its repayment and exit decisions. We integrate over f consistently—the same draw determines the transfer amount, the survival decision, and the next-period obligations—by quadrature over the Pareto distribution within each cell of the augmented state space.

Transfer amounts are then discretized into $K = 50$ bins (a finer partition leaves the results essentially unchanged). For the grant, the bin only matters in the program period, so the transition solver is the baseline one with a modified first period. For the guaranteed loan, the bin is an additional state variable for the following thirty years: each bin carries a fixed mortgage-style annual payment, computed at a rate of 2.5 percentage points above the risk-free rate, that enters the firm's budget constraint; a firm that defaults sheds its government loan along with its private debt, and the government absorbs the loss. The backward pass therefore solves $K + 1$ coupled value functions per period (parallelized across bins), and the government's net fiscal cost is the initial disbursement minus the present value of collected payments. The welfare effects of the programs are computed from the equilibrium consumption path as one-year consumption-equivalent variations.

D Additional Figures

Figure A.1: Sensitivity-based identification of κ

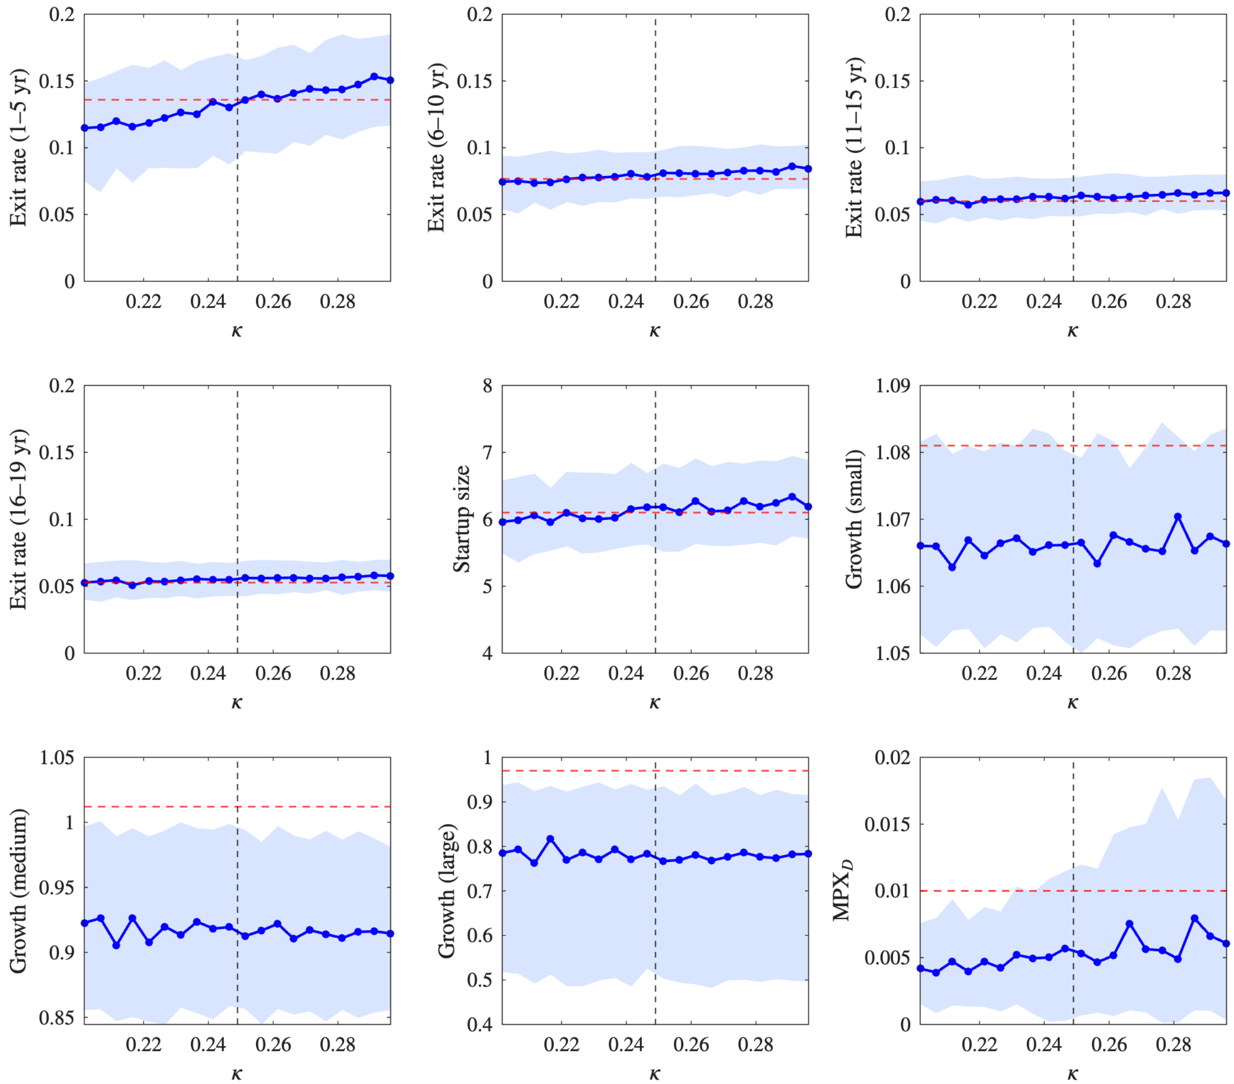

Notes: Each panel shows how one calibration moment (as reported in Table 4) varies with κ under a Sobol sensitivity design following the identification argument of Daruich (2026). All parameters are independently drawn from a 20% band around their estimated values. Solid lines denote the median across draws and shaded areas the interquartile range.

Figure A.2: Ergodic Distribution of Firms by Permanent Productivity Type

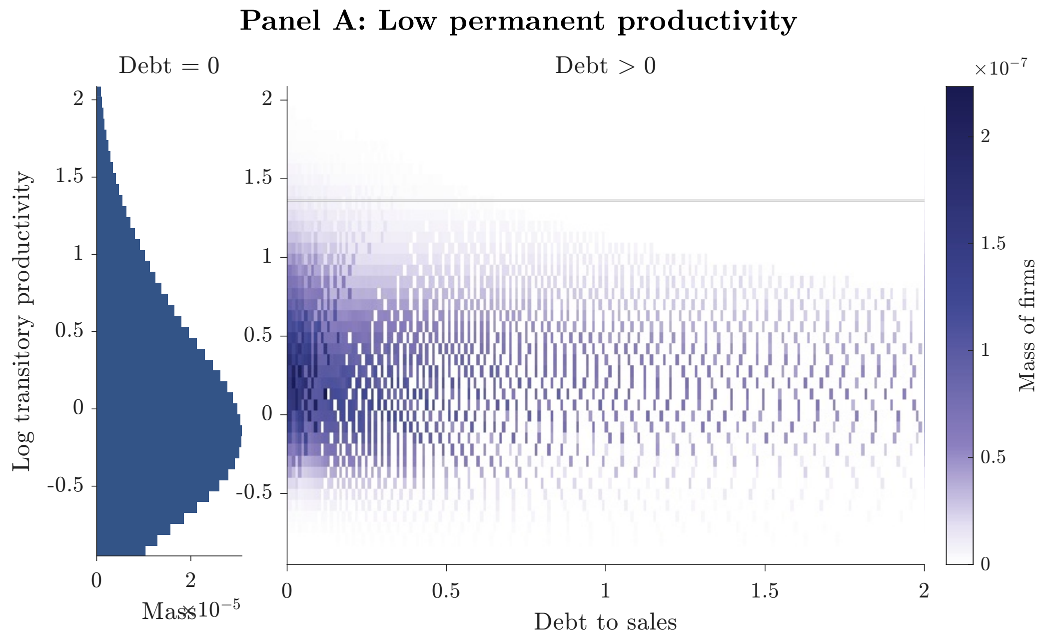

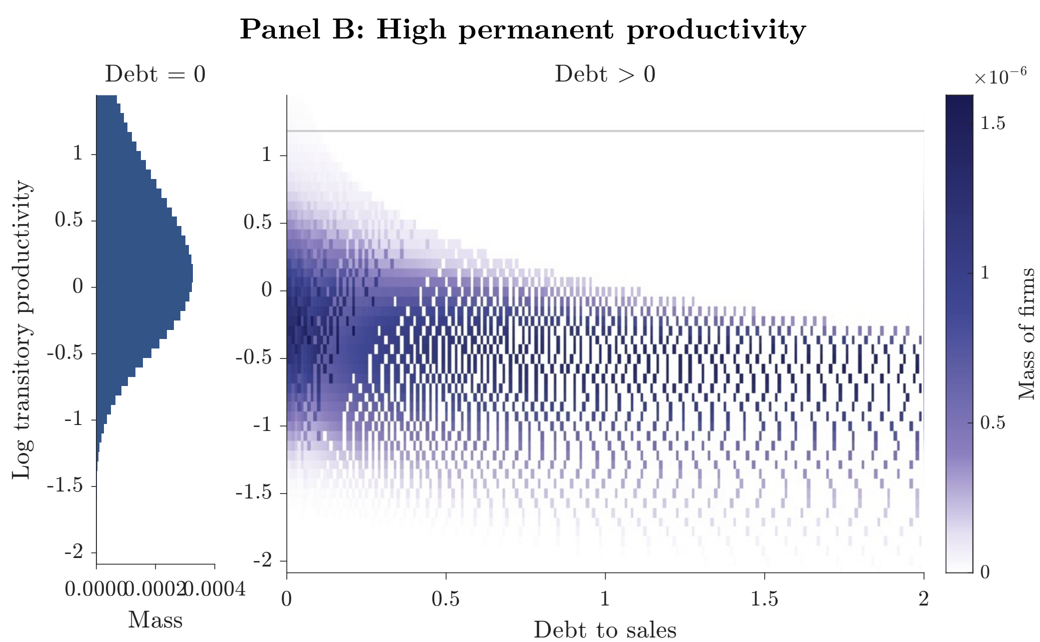

Notes: This figure displays the stationary distribution of firms over transitory log productivity and the ratio of initial debt obligations to sales for low (Panel A) and high (Panel B) permanent productivity types. Within each panel, the left subplot shows firms with zero debt; the right subplot shows firms with positive debt, where color indicates mass.

Figure A.3: Exit and default patterns: low permanent productivity

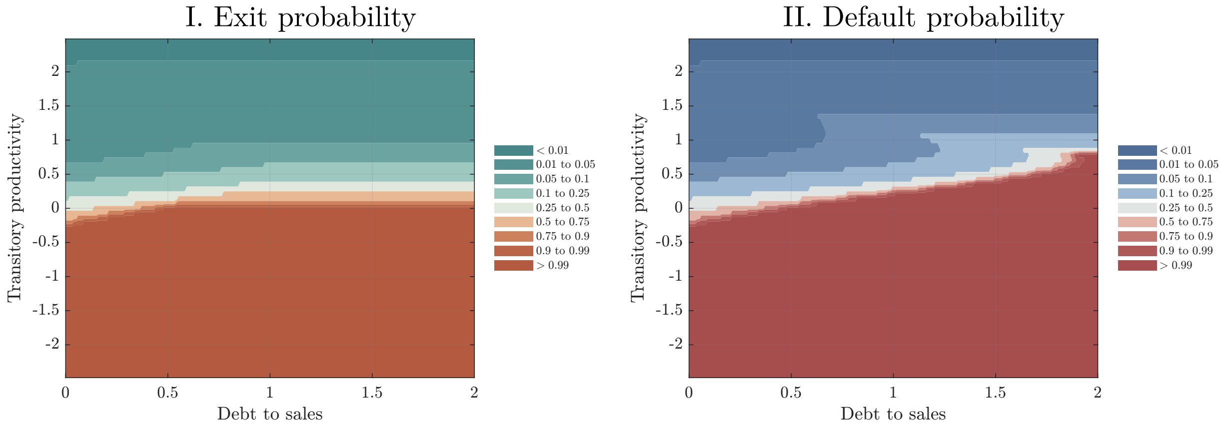

Notes: This figure presents the exit (left panel) and default (right panel) probabilities as a function of start-of-period transitory productivity (y-axis) and debt-to-sales (x-axis), for a low permanent productivity firm. Higher transitory productivity is associated with lower exit and default probabilities, while higher debt-to-sales is associated with higher exit and default probabilities.

Figure A.4: Exit and default patterns: high permanent productivity

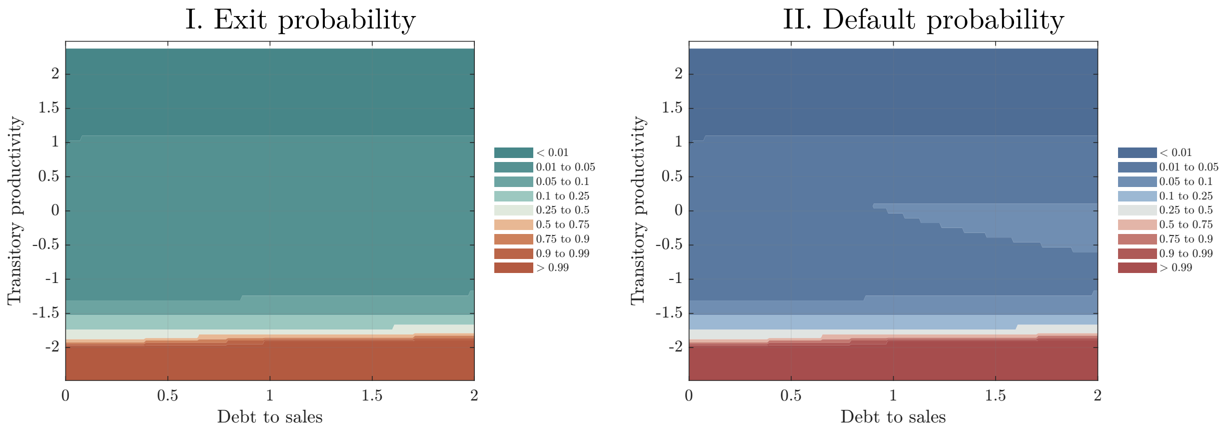

Notes: This figure presents the exit (left panel) and default (right panel) probabilities as a function of start-of-period transitory productivity (y-axis) and debt-to-sales (x-axis), for a high permanent productivity firm. Higher transitory productivity is associated with lower exit and default probabilities, while higher debt-to-sales is associated with higher exit and default probabilities.

E Additional Tables

Table A.1: Characteristics of Exiters vs. Incumbents

	Incumbents	**Exiting firms**	**T-statistic**
Age	16.92	14.92	7.89***
Sales	17.41	11.42	13.74***
Sales growth	12.76	5.27	3.13***
Capital stock	3.75	2.52	11.35***
Employment	52.96	46.57	2.21**
Labor productivity	1.65	−9.64	2.44**

Notes: This table compares firm fundamentals between incumbent firms and exiters. Sales and capital are expressed in millions of constant 2017 US dollars. Sales growth is the year-over-year percent change in sales. Employment corresponds to the number of employees. Labor productivity is computed as the deviation from the year-industry average, as output in millions of constant 2017 dollars per employee. The reported t-statistic tests for equality of means between exiters and incumbents. Statistical significance at the 10, 5, and 1 percent levels is denoted by *, **, and ***, respectively.

Table A.2: Propensity to Exit and Debt-to-Assets

	(1)	(2)	(3)	(4)
debt ratio	0.023***	0.020***	0.020***	0.018***
	(0.001)	(0.002)	(0.004)	(0.004)
Observations	119K	119K	24K	16K
R^2	0.007	0.009	0.009	0.010
Industry-Year FE	Yes	Yes	Yes	Yes
Firm Fundamentals				
Main controls	No	Yes	Yes	Yes
Additional	No	No	No	Yes
Sample	Full	Full	Census	Census

Notes: The dependent variable is a firm-level exit indicator equal to one if the firm exits the market in year t. The debt ratio is defined as debt to assets in year $t-1$. Main firm controls include a young firm indicator (five years or younger), firm size fixed effects (sales quintiles at the sector–year level), sales growth, and (log) capital stock; additional controls further account for (log) employment and (log) labor productivity. Columns (1) and (2) use the full sample while columns (3) and (4) use the Census sample. All regressions include industry–year fixed effects. Standard errors (in parentheses) are clustered at the industry–year level. Statistical significance at the 10, 5, and 1 percent levels is denoted by *, **, and ***, respectively.

Table A.3: Orbis Sample Coverage Relative to EU KLEMS — Italy (%)

Year	Employment	Wage bill	Output
2002	67	72	81
2003	59	91	88
2004	49	74	79
2005	50	76	81
2006	61	79	83
2007	62	79	84
2008	63	81	85
2009	60	82	88
2010	61	85	88
2011	76	86	91
2012	77	85	91
2013	79	87	91
2014	80	89	93
2015	82	89	95

Notes: Each cell reports the ratio (in %) of the Orbis sample aggregate to the corresponding EU KLEMS national account figure for the Italian manufacturing sector.

Table A.4: Firm-Size Distribution: Orbis vs. Eurostat — Italy (2012)

	Orbis		Eurostat	
Size class	Empl.	Output	Empl.	Output
<10 employees	.09	.06	.15	.09
10–19 employees	.12	.09	.15	.09
20–49 employees	.19	.14	.18	.14
50–249 employees	.29	.27	.25	.26
250+ employees	.31	.44	.27	.42

Notes: Each cell reports the share of total employment or total output accounted for by each size class in the Orbis sample and in Eurostat's Structural Business Statistics (SBS), for the year 2012.

Table A.5: Summary Statistics — Italy

	Mean	St. dev.	p25	p75	N
Firm fundamentals					
Age	19.86	15.51	8.00	28.00	906,724
Sales	3.00	4.99	0.25	3.00	1,967,190
Sales growth (%)	2.12	68.08	−12.27	16.87	1,577,620
Capital stock	0.70	1.29	0.02	0.63	1,969,904
Employment	31.65	177.94	4.00	23.00	1,342,815
Employment growth (%)	0.81	36.69	−5.13	8.00	1,057,035
Labor productivity	0.29	30.72	0.09	0.26	1,342,809
Financial characteristics					
Assets	3.12	5.14	0.32	3.08	1,970,100
Cash	0.18	0.35	0.00	0.15	1,946,525
Accounts receivable	0.83	1.53	0.00	0.79	1,969,980
Non-current assets	1.98	3.30	0.17	1.99	1,946,429
Liabilities	2.20	3.52	0.24	2.24	1,970,084
Short-term bank debt	0.35	0.78	0.00	0.23	1,970,084
Accounts payable	0.64	1.17	0.00	0.61	1,970,100
Non-current liabilities	1.07	1.60	0.12	1.20	1,970,084

Notes: Sales, capital stock (book value of tangible fixed assets), and all financial variables are expressed in millions of constant euros (2010 base year) and are winsorized at the 5–95% level each year. Employment corresponds to the number of employees. Labor productivity is the ratio of real operating revenue (millions of 2010 euros) to employment.