

Cyclical Fluctuations, Financial Frictions, and Productivity Differences across Firms*

Luca Guerrieri Jinill Kim Arsenii Mishin
Federal Reserve Board Korea University HSE University
Bank of Korea

June 9, 2026

Abstract

We bridge two literatures that have largely developed in parallel: the endogenous growth literature, which uses tractable representative-agent models to study the evolution of aggregate TFP, and the misallocation literature, which studies why financial frictions prevent capital from flowing to its most productive uses across heterogeneous firms. We build a tractable model that imports key features of the misallocation literature into a representative-agent framework. Credit rationing arising from adverse selection and moral hazard prevents capital from flowing efficiently to the most productive firms. As credit rationing becomes more or less intense, aggregate TFP acquires an endogenous component, accounting for about 30 percent of the variance of TFP at business cycle frequencies. We show that our tractable model can match key features of the observed distribution of productivity across firms and its co-movement with output growth and credit conditions in the data.

JEL Classification: E23, E32, E44

Keywords: productivity dispersion, endogenous productivity, financial intermediation, business cycles

*The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System, of any other person associated with the Federal Reserve System, of the Bank of Korea, or of any other person associated with the Bank of Korea. An online appendix is available at the following link: https://www.lguerrieri.com/downloads/Appendix_GKM.pdf

1 Introduction

Two strands of the macroeconomics literature have developed largely in parallel without much cross-fertilization. The first is the endogenous growth literature, which uses representative-agent frameworks to study the low-frequency evolution of aggregate total factor productivity (TFP) based on foundational contributions by Lucas (1988), Romer (1990), and Aghion and Howitt (1992).¹ More recently, Comin and Gertler (2006) showed how endogenous technology adoption and R&D generate medium-frequency TFP dynamics that persist well beyond the business cycle. These models are tractable enough to be estimated and embedded in larger quantitative frameworks, but they abstract from the cross-sectional distribution of productivity across firms.²

The second is the misallocation literature, which studies why capital and labor are not efficiently allocated across heterogeneous firms and what this misallocation implies for aggregate productivity. A large body of work documents that productivity dispersion across firms within narrowly defined industries is substantial and persistent, and shows that financial frictions are a leading culprit: credit constraints, information asymmetries, and moral hazard prevent resources from flowing to their most productive uses.³ The theoretical models that can account for these patterns in full generality are necessarily complex—they require tracking distributions of firm-level state variables and are difficult to embed in standard representative-agent business cycle frameworks. Yet bridging these literatures has direct macroeconomic stakes. The aggregate Solow residual is not a primitive of the economy: it aggregates up from the allocation of inputs across firms with different productivity levels, so when financial frictions cause that allocation to vary with credit conditions, measured TFP inherits a component that is endogenous to macroeconomic developments rather than

¹Lucas (1988), Romer (1990), and Aghion and Howitt (1992) established how human capital accumulation, R&D, and innovation endogenously determine the level and growth rate of productivity. King and Rebelo (1990) and Rebelo (1991) embedded these mechanisms in calibrated representative-agent frameworks, showing how endogenous growth alters the predictions of RBC models with implications for policy analysis.

²For further examples in this tradition, see Jones et al. (2005) and Dang et al. (2012) on endogenous growth with business cycle implications, and Gornemann et al. (2025) for a quantitative open-economy application with endogenous productivity.

³Foundational contributions include Foster et al. (2008), Restuccia and Rogerson (2008), Hsieh and Klenow (2009), Midrigan and Xu (2014), Moll (2014), Gopinath et al. (2017), and David and Venkateswaran (2019).

reflecting technological change alone.

This paper bridges these two literatures. Our central contribution is to show how key features of the misallocation literature, such as endogenous productivity dispersion driven by financial frictions, can be imported into a tractable representative-agent setting. In our model, each firm is endowed with two technologies: a production technology, whose idiosyncratic efficiency is private information, and a financial intermediation technology, which is available to all firms on equal terms. Because firm-level production efficiency is not directly observed by investors and the claims of firms are not credible, this informational friction results in all firms being financed with an aliquot share of household savings, regardless of productivity. In a secondary market for funds, where the intermediation technology of a firm allows a little more screening of productivity to be performed, firms compare the returns from their two technologies and sort themselves into lenders, firms deploying their intermediation technology; borrowers that default strategically without producing; and producers. This sorting depends on the realization of the production-technology draw, and all we need to do is characterize the productivity points that demarcate different groups of firms. As these cutoff points respond to aggregate macroeconomic conditions, the resulting model delivers endogenous dynamics for aggregate TFP akin to those in the endogenous growth literature: Roughly 30 percent of the variance of TFP at business cycle frequencies is endogenous in our model.

But our contribution extends beyond tractability. The framework captures strategic default as an equilibrium outcome, not a knife-edge case ruled out by contract design, and the equilibrium variation in the mass of defaulting firms is quantitatively important: it accounts for roughly one-third of the endogenous TFP component, over and above what a single-cutoff model without strategic default can generate. Strategic default thus serves a dual purpose: it amplifies the endogenous TFP component by concentrating production among high-productivity firms, and in our simulations also stabilizes the dynamic path of that component—when aggregate conditions shift, firms near the default threshold adjust first, distributing the reallocation burden across two margins rather than concentrating it in a single threshold, which yields a smoother and more persistent endogenous TFP response than the no-default variant produces.

Our approach avoids a classic criticism of the endogenous growth literature. Models of endogenous growth can lead to strong predictions and policy recommendations based on long-run growth effects that are difficult to ascertain empirically. In our model, the endogenous growth effects are not permanent. The tractable structure allows us to deploy standard methods for an empirical verification. While the endogenous growth literature has generated deep insights about long-run forces shaping aggregate TFP, empirically isolating those forces has proven difficult: Aggregate time series data cannot readily distinguish neoclassical from endogenous growth models (Pack, 1994), and the low-frequency statistical models underlying many empirical exercises are nearly indistinguishable from one another at typical macroeconomic sample lengths (Müller & Watson, 2008). By contrast, we can use data on the dispersion of productivity for firms within narrowly defined industries in support of the micro-foundations of our model.

As in Stiglitz and Weiss (1981), SW henceforth, credit rationing is at the center of our model, but it arises from some additional incentives built into our model. In SW, default arises involuntarily when project returns fall short of the promised repayment; the interest rate affects bank returns by changing the composition of applicants (adverse selection) and the riskiness of projects chosen (incentive effect/moral hazard). Our model adds a different channel: firms with sufficiently low productivity find it profitable to divert borrowed funds to an outside option rather than produce, making default a strategic, equilibrium choice.⁴ This moral hazard margin, different from the one in SW, generates a two-tier sorting structure. In our model interest rates influence both the strategic default rate and the average productivity of firms that choose to borrow, produce, and repay their loans.⁵

In our model, misallocation is not static. As aggregate conditions change, the thresholds governing which firms lend, borrow, and default fluctuate, and with them the dispersion of productivity across firms and aggregate TFP. The model thus encompasses two sources of

⁴There are other important differences between our analysis and the one in Stiglitz and Weiss (1981). We develop a general equilibrium model whereas theirs is a static partial equilibrium model. The model in Stiglitz and Weiss (1981) has no capital accumulation, no goods market, no aggregate TFP, and no business cycle dynamics. The bank's interest rate is a parameter, not an endogenous price that co-moves with output, investment, and consumption as in our model.

⁵In an entrepreneurship model with multiple sources of financial frictions, Paulson et al. (2006) showed—via structural estimation on Thai data—that moral hazard due to information asymmetry is the dominant source.

aggregate TFP: a standard exogenous component and an endogenous component driven by changes in the distribution of productive activity across firms. Under our baseline calibration, this endogenous channel generates meaningful differences in impulse responses to TFP shocks relative to models without financial frictions and productivity dispersion. In the efficient limit of our model, all credit flows to the most productive firm and the economy collapses to a standard RBC model.

Despite its relative simplicity, the model is sufficiently flexible to match key features of the data. We calibrate the model to standard targets for credit market conditions and productivity dispersion, and find that it comes close to replicating several untargeted moments. In particular, OLS regressions estimated on model-simulated data reproduce the signs of their counterparts in U.S. data: within-industry productivity dispersion narrows when GDP growth runs above average, and delinquencies on business loans rise when growth is depressed and when real interest rates are high.

2 Literature Review

Our work has many antecedents. It is related first to papers that showcase models with costly state verification, such as Carlstrom and Fuerst (1997), Bernanke et al. (1999), and Christiano et al. (2014). These papers established that the agency costs of financial intermediation can substantially amplify and propagate aggregate shocks, a foundational insight that allows us to go further. A common modelling device for these papers is the assumption that credit needs to be allocated to firms with heterogeneous productivity. However, the productivity dispersion is static and exogenously given—it cannot respond to economic conditions because the distribution of idiosyncratic productivity is assumed fixed and independent of aggregate developments. Default in those models is also not strategic: it stems from the inability to repay after uncertainty resolves unfavorably. A further distinction concerns the source of heterogeneity: in those frameworks it is idiosyncratic project-return or entrepreneurial risk, rather than the producing-firm productivity-selection margin that is central to our model. In Christiano et al. (2014) in particular, idiosyncratic risk varies over time, but exogenously, rather than responding to the equilibrium sorting of firms. By

contrast, in our model default is strategic, varies with economic conditions, and feeds back into productivity dispersion and aggregate TFP.

Our work is also related to papers that include endogenously evolving productivity dispersion. Khan and Thomas (2013) were the first to explore the endogenous TFP channel in a quantitative DSGE setting where reallocation of capital across heterogeneous firms determines aggregate TFP. Their model works through the interaction of two frictions—collateralized borrowing constraints and partial capital irreversibility—that together distort capital allocation and generate persistent endogenous TFP movements. Our model captures the same channel through a more parsimonious mechanism: strategic default sorts firms by productivity and generates credit rationing without requiring capital-level heterogeneity or (S,s) adjustment rules. And unlike their model, which is designed to rule out default in equilibrium, default is an equilibrium outcome in ours. Khan and Thomas’s rich framework allows them to match detailed microeconomic moments on firm investment behavior that our model does not target; the tradeoff is that their solution algorithm requires tracking a three-dimensional firm distribution. By contrast, the simplicity of our framework allows us to embed it in many variants of an RBC model, and we view it as a building block that others can readily reuse or extend.

Our emphasis on tractability connects us to Buera and Moll (2015), who developed a model where heterogeneous firms face collateral constraints and shocks to those constraints open productivity wedges. In their benchmark model, idiosyncratic firm productivity is independently and identically distributed (i.i.d.) across entrepreneurs and over time, so the cutoff governing which firms produce is determined exogenously by the collateral constraint rather than by equilibrium sorting. When they allow for persistent idiosyncratic shocks—treated in their online appendix under logarithmic utility to preserve tractability—the conclusions were the same as in the i.i.d. case. Under logarithmic preferences, moreover, a tightening of collateral constraints is exactly isomorphic to an exogenous TFP shock, and no other shocks are considered; the endogenous TFP component therefore cannot be separately identified from the exogenous one. Our model delivers an endogenous TFP component even in the i.i.d. limit, and the strategic-default margin generates an additional channel that does not arise in their framework.

Our work is also related to Liu and Wang (2014) and Dong et al. (2025), who studied heterogeneous firms under collateral constraints in similar settings. In Liu and Wang (2014), borrowing constraints are governed by an exogenous parameter rather than by the equilibrium sorting of firms across lending, production, and default. By design, no firm defaults in equilibrium in their framework; default is instead an equilibrium outcome in ours, and it is precisely the endogenous variation in the mass of defaulting firms that drives an important part of the endogenous TFP response. A further difference is that the endogenous TFP channel in our model operates through equilibrium sorting at the default margin, without relying on entry and exit costs. Dong et al. (2025) is a targeted extension of Liu and Wang (2014) that focuses on the effects of turbulence shocks. In their model, misallocation arises from the joint presence of credit constraints and production distortions; turbulence operates by tightening the borrowing constraints of high-productivity firms and amplifying pre-existing misallocation. In our model, misallocation arises endogenously from financial frictions alone, without requiring pre-existing production distortions.

The misallocation literature more broadly provides important context for our work. Restuccia and Rogerson (2008) and Hsieh and Klenow (2009) established that policy distortions and wedges between marginal products account for large aggregate productivity losses. Midrigan and Xu (2014) and Moll (2014) emphasized financial frictions as a central driver of misallocation and productivity losses. Moll (2014) established an important benchmark: With serially persistent idiosyncratic productivity, entrepreneurs can gradually self-finance out of credit constraints, which reduces misallocation and allows TFP to recover over time. Without serial correlation, however, that particular self-financing channel disappears and TFP is fully driven by exogenous factors. Our paper goes further by demonstrating that endogenous TFP dynamics arise generically even when idiosyncratic productivity is i.i.d. First, an equilibrium-selection channel makes the cutoffs for lending, borrowing, and default responsive to aggregate conditions, producing endogenous TFP variation independent of wealth dynamics. Second, explicitly modeling strategic default creates an additional, quantitatively relevant source of TFP fluctuations through firms' default choices and their general-equilibrium effects. Together these mechanisms substantially broaden the range of financial-contracting features that can drive aggregate misallocation and TFP dynamics be-

yond the scope of Moll (2014). David and Venkateswaran (2019) and Gopinath et al. (2017) documented the importance of capital misallocation both in the United States and in Southern Europe. On the firm dynamics side, Clementi and Hopenhayn (2006) and Cooley and Quadrini (2001) developed theories of financing constraints and firm dynamics that inform our modeling of credit frictions.

Our paper also contributes to a body of work exploring heterogeneous firms, financial frictions, and default. Gilchrist et al. (2014) studied credit spreads under uncertainty shocks in a model with default. Arellano et al. (2019) examined the role of uncertainty shocks in a model with noncontingent debt and equilibrium default. Gomes and Schmid (2021) developed a model with endogenous default where firms vary in leverage and studied the implications for credit spreads.⁶

Finally, our paper is related to an extensive literature explaining productivity differences across firms. Beyond financial frictions, alternative explanations include slow learning of average productivity in the face of noise shocks, as in Jovanovic (1982); the coexistence of older and newer technology vintages, as in Caballero and Hammour (1994), Caballero and Hammour (1996), and Caballero and Hammour (1998); search and matching frictions between workers and firms, as in Barlevy (2002); and imperfect product substitutability, as in Bernard et al. (2003) and Melitz (2003).

3 Data Motivation

We document that within-industry productivity differences in the U.S. manufacturing sector are sizable and pervasive, and that they tend to widen when GDP growth runs below average. The summary statistics constructed here are the same ones that we use to assess our theoretical model.

We draw dispersion in productivity across 86 four-digit NAICS manufacturing industries from experimental productivity dispersion statistics, Dispersion Statistics on Productivity

⁶Jo et al. (2021) and Ottonello and Winberry (2020) also introduced financial frictions and firm default, but our model is more tractable and intuitive by incorporating moral hazard and strategic default; they focused on size and leverage distributions, whereas our model emphasizes consistency with productivity dispersion. Guo et al. (2025) discussed business cycles with firms' private information and equity financing.

(DiSP), derived from Census Bureau microdata.⁷ The most recent release of DiSP covers the years 1987–2021 on an annual basis. Our analysis focuses on the second-moment measure of establishment-level total factor productivity.

Figure 1 shows the evolution of TFP dispersion over time. The top panel considers the percentage difference in productivity between the establishment at the 90th percentile of the TFP distribution and the 10th percentile establishment in the same industry. The solid line denotes the average of these within-industry 90–10 gaps, weighted by each industry’s contribution to GDP.⁸ The units on the vertical axis are percentage points: a reading of one hundred means that the plant at the 90th percentile makes twice as much output as the 10th percentile plant with the same measured inputs. By that yardstick, the average 90–10 gap is sizable throughout the sample, fluctuating around 100 percent without an evident trend.

The dot-dashed bands describe how that gap varies across industries. Year by year, we rank the 86 industries by their own 90–10 gap and plot the gap of the industry at the 15th percentile and at the 85th percentile of the cross-industry distribution. The band is fairly symmetric around the solid line, and even the industry at the 15th percentile across industries still has a sizable 90–10 gap. The average TFP gap is therefore not driven by a handful of outliers: large within-industry productivity differences are pervasive across the manufacturing sector.

For robustness, the bottom panel repeats the same analysis for the interquartile range (IQR). Within-industry productivity differences remain important quantitatively when we focus on segments closer to the center of the distribution, not just the tails.

Turning to cyclical variation, the shaded areas in Figure 1 mark recessions as dated by the National Bureau of Economic Research. Reading the solid line in the top panel against those shaded windows—and against the lower-frequency variation in GDP growth that they highlight—the 90–10 gap tends to widen in the years when GDP growth runs below its sample average and, most visibly, during NBER recessions. The same broad pattern is apparent in the bottom panel for the interquartile range.

⁷This dataset is described in Cunningham et al. (2023).

⁸The industry value-added data used to construct the weighted average TFP dispersion shown in Figure 1 is from the Bureau of Economic Analysis. To construct the industry weights, we aggregate the value added of multiple NAICS manufacturing industries that correspond to the same BLS code.

4 Model

Our analysis focuses on the key role of credit markets in reallocating capital among heterogeneous firms. Each firm in the goods sector is endowed with two technologies. The first is a production technology that combines capital and labor using an idiosyncratic efficiency parameter ω drawn from a distribution on the unit interval; this draw is transitory and privately observed by the firm. The second is a financial intermediation technology, available on equal terms to every firm, which allows it to lend funds to other firms at the prevailing inter-firm rate. Firms also experience persistent aggregate technology shocks that shift the productivity of the production technology for all firms simultaneously. Because the idiosyncratic efficiency of each firm’s production technology is private information, households allocate an aliquot share of their savings to each firm without conditioning on ω . We will show that firms with sufficiently high realizations of ω borrow from firms with low realizations to purchase additional capital, while the latter find it optimal to deploy their financial intermediation technology instead of producing.

Unlike in a frictionless credit market, two distinct informational frictions interact to generate misallocation of capital and reduced productivity in our framework. The first is adverse selection: the production-technology efficiency ω above some threshold level is private and unobservable to outside investors, so the composition of borrowers cannot be screened perfectly. Because all eligible firms could credibly claim high productivity, offering to pay a higher rate to borrow more funds lacks credibility, and credit is rationed as in Stiglitz and Weiss (1981). The second friction is moral hazard: firms with sufficiently low realizations of ω find it profitable to divert borrowed funds to an outside option rather than produce and repay, making default a strategic, equilibrium choice. It is this moral hazard channel—not simply the adverse-selection problem—that limits the borrowing capacity of even the most efficient firm, because lenders must price in the possibility that some borrowers will divert.

Our interpretation of inter-firm lending reflects the broader role of credit markets in allocating credit. As Bernanke and Gertler (1990) note, this lending can be viewed as intermediated by competitive financial institutions that neither use resources nor earn profits

in equilibrium. Variation in the cutoff points we just discussed will also lead to endogenous changes in the size of the intermediation sector, as the boundary between firms deploying their production technology and those deploying their financial intermediation technology shifts with aggregate conditions.

Aside from the endogenous financial intermediation, the other main innovation of the model consists of tracking the dispersion of productivity across goods-producing firms in a tractable way. Relative to a standard model, the complications will consist of characterizing three additional equilibrium objects. The first two objects are a couple of cutoff points within the distribution of firm productivity. A lower cutoff point will separate firms with such low productivity that lending to more productive firms will be more profitable than producing. Just on the other side of this cutoff will be more-productive firms that borrow in the inter-firm market. An upper cutoff point will demarcate firms that borrow and default from those that do not default. The third additional object is the price at which funds are lent across firms.

As these cutoff points and the cost of credit fluctuate, productivity differences across firms will also fluctuate, as will default rates, together with the overall efficiency of the economy. We will show that despite the simplicity of our model, key characteristics of these fluctuations are well aligned with U.S. data.

Firms are at the center of our model, but we start with households for ease of exposition.

4.1 Households

There is an infinitely-lived representative household that has preferences over consumption and labor, respectively C_t and H_t , in line with King et al. (1988). The household solves the following problem:

$$\max_{\{A_{t+\tau}, C_{t+\tau}, H_{t+\tau}, B_{t+\tau}^H\}_{\tau=0}^{\infty}} E_t \sum_{\tau=0}^{\infty} \beta^{\tau} \left[\ln(C_{t+\tau} - \nu_{ct+\tau}) - \frac{\vartheta}{1+\nu} (H_{t+\tau})^{1+\nu} \right], \quad (1)$$

subject to

$$C_{t+\tau} + A_{t+\tau} + B_{t+\tau}^H = R_{t+\tau}^A A_{t+\tau-1} + W_{t+\tau} H_{t+\tau} + R_{t+\tau-1}^B B_{t+\tau-1}^H + \Pi_{t+\tau} + T_{t+\tau} + \Xi_{t+\tau}. \quad (2)$$

In period $t + \tau$, the household chooses how much to consume, $C_{t+\tau}$, and work, $H_{t+\tau}$. The term $\nu_{ct+\tau}$ is an exogenous shock to consumption, which follows an autoregressive process of order one:

$$\nu_{ct} = \rho_\nu \nu_{ct-1} + \varepsilon_{\nu t}, \quad \varepsilon_{\nu t} \sim \mathcal{N}(0, \sigma_\nu^2). \quad (3)$$

The parameter β discounts future utility. The parameter ϑ captures the disutility from working, and the parameter ν governs the Frisch elasticity of labor supply. The household also chooses the amount of assets $A_{t+\tau}$ and a government bond $B_{t+\tau}^H$. The household enters period $t + \tau$ with assets $A_{t+\tau-1}$ which carry a state-contingent return $R_{t+\tau}^A$ from equity ownership of the goods-producing firms (which act as both producers and financial intermediaries) and receives a non-state contingent return $R_{t+\tau-1}^B$ from its holdings of the government bond $B_{t+\tau-1}^H$. It also supplies $H_{t+\tau}$ units of labor at the wage rate $W_{t+\tau}$. The term $\Pi_{t+\tau}$ captures profits from ownership of capital-producing firms; goods-producing firms earn zero profits in equilibrium due to perfect competition, while capital-producing firms can earn non-zero profits in equilibrium. The term $T_{t+\tau}$ represents a lump-sum transfer from the government, which runs a balanced budget, period by period (defined in equation (29) of Section 4.4). Finally, the term Ξ_t , defined in equation (30) of Section 4.4, refers to transfers that the household receives because firms that take the outside option, which is described below, are subject to a haircut on the returns from the outside option.

The first-order conditions for assets, consumption, labor, and government bonds are derived in Section C.1 of the [online appendix](#); see equations (C.1)–(C.4), where λ_{ct} is the Lagrange multiplier attached to the budget constraint (2).

4.2 Production and Financial Intermediation

Firms in the goods sector are at the heart of our model. The firms in this sector follow a two-period overlapping structure from period t to $t + 1$. This device is familiar from decentralizations of the RBC model that rely on equity contracts to allocate household savings to firms. Table 1 provides a roadmap to the sequence of choices made and actions taken by firms within each period. We highlight with a **boldface** font the choices or actions that are specific to our model. Lowercase letters will denote variables of individual firms; we

shall reserve uppercase letters for aggregate variables.

Firms operate in a perfectly competitive market, producing a homogeneous good. Ex ante, each firm is endowed with two technologies: a production technology and a financial intermediation technology. An idiosyncratic efficiency parameter $\omega \in [0, 1]$, drawn from an identical distribution at the start of each period, governs how productive the firm would be if it chose to deploy its production technology. The financial intermediation technology, which consists of lending funds in the inter-firm market, is identical across firms. Depending on the realization of ω , firms endogenously choose which technology to deploy: those with high ω find production more profitable and borrow to expand their capital stock, while those with low ω find it optimal to act as financial intermediaries and lend their funds to higher-productivity firms.

In a slight abuse of notation, the term ω will do double duty by denoting the idiosyncratic production efficiency and by serving as an index for other firm-specific variables. Since households have no visibility into the idiosyncratic production efficiency of firms, each firm receives an aliquot share of the households' savings, a_t . So, notice that a_t does not depend on ω in equilibrium.

Starting from the production function, the output of an individual firm in period $t + 1$, $y_{t+1}(\omega)$, is governed by

$$y_{t+1}(\omega) = \omega Z_{t+1} k_t(\omega)^\alpha h_{t+1}(\omega)^{1-\alpha}. \quad (4)$$

The terms $k_t(\omega)$ and $h_{t+1}(\omega)$ denote the levels of capital and labor inputs used by the firm. The idiosyncratic productivity ω follows the cumulative distribution function $\mu(\omega)$ on the interval $[0, 1]$, satisfying $\mu(0) = 0$, $\mu(1) = 1$, and $\mu'(\omega) > 0$. The aggregate technology shock Z_{t+1} evolves according to

$$\log Z_{t+1} = \rho_z \log Z_t + \varepsilon_{t+1}^z, \quad \varepsilon_{zt} \sim \mathcal{N}(0, \sigma_z^2), \quad (5)$$

where the parameter ρ_z governs the persistence of the shock process.

Firms make plans in period t to produce in period $t + 1$. After their idiosyncratic productivity ω is known, the intermediation market opens. Depending on their productivity, some firms borrow while others lend at the predetermined rate ρ_t . Firms that borrow decide

whether to use an outside option by purchasing the government bond or produce.⁹ If they choose to produce, they purchase physical capital.

Firms can walk away from loans and default. If a firm decides to default, it can retain a fraction $\Theta_t(\omega)$ of the funds borrowed from other firms $b_t(\omega)$. We assume that $\Theta_t(\omega)$ increases in ω and is concave. Section B of the [online appendix](#) provides a structural interpretation for our modeling of $\Theta_t(\omega)$. For simplicity, we make the average fraction of funds that can be diverted equal to θ , so

$$\theta = \int_{\bar{\omega}_t}^{\bar{\bar{\omega}}_t} \Theta_t(\omega) \frac{\mu'(\omega)}{\mu(\bar{\bar{\omega}}_t) - \mu(\bar{\omega}_t)} d\omega, \quad (6)$$

where, as will be shown later, $\bar{\omega}_t$ and $\bar{\bar{\omega}}_t$ define the boundaries of defaulting firms, and $\frac{\mu'(\omega)}{\mu(\bar{\bar{\omega}}_t) - \mu(\bar{\omega}_t)}$ gives the corresponding probability density over that interval.¹⁰

Creditor firms are assumed not to be able to go after any funds set aside for the outside option. The diverted assets are placed with the government bond; taking this outside option incurs a haircut ξ on the bond return, which is rebated in lump-sum fashion to households. The precise outside-option payoff and the role of ξ are set out below when we characterize the optimization problem of defaulting firms.

Whereas the idiosyncratic productivity of firms is not directly observed by households, we assume that firms have a bit more information. While still far from perfect, a screening technology allows lending firms to tell whether borrowers can be expected to make more than the outside option by producing. This technology prevents firms with the lowest levels of production efficiency from borrowing and defaulting, thereby supporting a mass of lenders among firms with $\omega \leq \bar{\omega}_t$.¹¹ Firms whose private productivity ω falls between $\bar{\omega}_t$ and $\bar{\bar{\omega}}_t$ find it advantageous to divert all available funds towards the outside option. Only firms that can

⁹The inter-firm loan rate ρ_t is pinned in equilibrium by the zero-profit condition of lending firms, equation (32): lenders must be indifferent between lending at rate ρ_t and taking the outside option themselves. In other words, ρ_t is not an exogenous constant but an endogenous price determined jointly with the two cutoffs by the equilibrium conditions characterized in Section 5.1.

¹⁰In Section C.5 of the [online appendix](#), we show that, as long as $\mu'(\bar{\omega}_t) \neq 0$ (a condition satisfied for our baseline calibration), the integral in equation (6) exists in the limiting case in which $\bar{\bar{\omega}}_t$ tends to $\bar{\omega}_t$.

¹¹A fully accurate screening technology would deliver the first-best allocation with only the most productive firm producing. The imperfect screening we assume guarantees instead that low-efficiency firms choose to lend rather than switch to borrowing and defaulting. Formally, the cutoff $\bar{\omega}_t$ is not set directly by the screening technology but is determined endogenously by the equilibrium indifference condition (31) and the zero-profit condition (32); the screening technology's role is to support this equilibrium by preventing deviations in which low- ω firms would mimic borrowers.

make higher returns by producing will produce.

With constant returns to scale production, the economy would attain the first-best allocation if the most efficient firm, the one with $\omega = 1$, could borrow all financing from other firms, allowing it to be the only firm to produce. This allocation is hindered by the presence of moral hazard due to asymmetric information. By limiting the borrowing capacity of the most efficient firm, moral hazard gives less efficient firms room to borrow from financial intermediaries. Asymmetric information refers to uncertainty about the quality of borrowers. The productivity level above $\bar{\omega}_t$ is private information that is not observed by other firms. Producing firms may seek more funds by offering to borrow above the prevailing lending rate. However this offer lacks credibility as all firms above a certain threshold could claim high productivity. The result is credit rationing, as in Stiglitz and Weiss (1981). Consequently, all inter-firm financial contracts are identical and do not depend on ω in equilibrium.

In the next three sections, we will consider three distinct optimization problems spanning the possible combinations of actions that firms can take. Firms can borrow from other firms and produce, borrow and default, or lend. We will show that firms sort themselves in each group depending on their idiosyncratic productivity, ω . For now, take this result as a hypothesis. The proof of the hypothesis is set up in Section 5 with details pushed to Section A of the [online appendix](#).

4.2.1 Firms Choosing to Produce ($\omega \geq \bar{\omega}_t$)

Firms acquire the capital stock required for production in period $t + 1$, by combining funds from households $a_t(\omega)$ with a loan $b_t(\omega)$ from the mass of lending firms; consequently,

$$b_t^{tot}(\omega) = a_t(\omega) + b_t(\omega), \tag{7}$$

where $b_t^{tot}(\omega)$ denotes the total borrowing of the firm. By contractual agreement, the funds borrowed can only be used to purchase capital; so,

$$b_t^{tot}(\omega) = Q_t k_t(\omega), \tag{8}$$

where Q_t is the price of capital in terms of consumption. Let $R_{t+1}(\omega)$ denote the rate of return on capital ownership and $\pi_{t+1}(\omega)$ describe the profit of the firm. The revenue of the firm includes the proceeds from the sale of output as well as from the sale of the undepreciated fraction of capital. The expenses encompass commitments associated with loan servicing as well as remuneration for labor services. Thus,

$$\pi_{t+1}(\omega) = z_{t+1}\omega k_t(\omega)^\alpha h_{t+1}(\omega)^{1-\alpha} + (1 - \delta)Q_{t+1}k_t(\omega) - R_{t+1}(\omega)b_t^{tot}(\omega) - W_{t+1}h_{t+1}(\omega), \quad (9)$$

where δ is the depreciation rate.

At time t , the problem of this producer of goods is to choose $b_t^{tot}(\omega)$, $k_t(\omega)$, and $h_{t+1}(\omega)$ to maximize expected profits in period $t+1$. Both capital and labor are committed in period t , before the firm-specific productivity ω is realized: capital is predetermined through credit rationing, and labor is hired one period in advance, reflecting the time required to contract, train, and match workers to firms. This maximization problem can be expressed as:

$$\max_{k_t(\omega), h_{t+1}(\omega), b_t^{tot}(\omega)} E_t \left\{ \beta \frac{\lambda_{ct+1}}{\lambda_{ct}} \pi_{t+1}(\omega) \right\}, \quad (10)$$

where $\beta \frac{\lambda_{ct+1}}{\lambda_{ct}}$ is the stochastic discount factor coming from the household's problem. This maximization is subject to (8).

Because both capital and labor are committed before ω is observed, and all producing firms are ex-ante identical, the optimal labor allocation is uniform across producing firms,

$$h_{t+1}(\omega) = \tilde{h}_{t+1}, \quad (11)$$

so the labor-capital ratio $h_{t+1}(\omega)/k_t(\omega) = \tilde{h}_{t+1}/\tilde{k}_t$ is equalized across producing firms (see Section C.3 of the [online appendix](#)). Aggregating over the producing segment yields the familiar aggregate labor-demand condition

$$W_t = (1 - \alpha) \frac{Y_t}{H_t}. \quad (12)$$

Using zero-profit pricing, the return on capital earned by a firm with productivity ω is

$$R_{t+1}(\omega) = \frac{\kappa_{t+1}}{Q_t} \omega + \frac{(1-\delta)}{Q_t} Q_{t+1}, \quad (13)$$

where $\kappa_{t+1} \equiv \alpha Z_{t+1} \left(\frac{H_{t+1}}{K_t} \right)^{1-\alpha}$ is a common factor across all producing firms. Because ω is strictly increasing, $R_{t+1}(\omega)$ is strictly increasing in ω : more productive firms earn higher returns. The factor $\frac{1}{Q_t}$ converts one unit of consumption into $\frac{1}{Q_t}$ units of capital; the term $\kappa_{t+1} \omega$ is the marginal product of capital at productivity ω ; and $(1-\delta)Q_{t+1}/Q_t$ captures the capital-gain component.

Upon observing ω , the producing firm determines the demand for inter-firm loans by solving the following problem:

$$\max_{b_t(\omega)|\omega} E_t \left[\beta \frac{\lambda_{ct+1}}{\lambda_{ct}} \left(R_{t+1}(\omega) (a_t(\omega) + b_t(\omega)) - R_{t+1}^A(\omega) a_t(\omega) - \rho_t b_t(\omega) \right) \right]. \quad (14)$$

It generates returns from owning capital by utilizing borrowed funds and distributes these returns to meet equity commitments of households and repay loans from other firms.

Note that, given constant returns to scale, a firm choosing to produce would not borrow if the cost of inter-firm funds exceeded the expected return from production. A sufficient condition for the corner solution in which such a firm borrows as much as possible is therefore $\rho_t \leq E_t R_{t+1}(\omega)$ for all $\omega \geq \bar{\omega}_t$; this inequality is verified in equilibrium for the calibrated model. Accordingly, supply conditions will have to determine how much these firms can borrow.

We use the zero-profit condition that holds under every state of nature to size the return paid to households by firms in this segment, i.e.

$$R_{t+1}^A(\omega) = \frac{R_{t+1}(\omega) (a_t(\omega) + b_t(\omega)) - \rho_t b_t(\omega)}{a_t(\omega)}. \quad (15)$$

4.2.2 Firms Choosing to Borrow and Default ($\bar{\omega}_t < \omega < \bar{\bar{\omega}}_t$)

The problem of the firm that diverts the borrowed funds by choosing the outside option can be described as follows:

$$\max_{b_t(\omega)|\omega} E_t \left[\beta \frac{\lambda_{ct+1}}{\lambda_{ct}} \left((R_t^B - \xi) (a_t(\omega) + \Theta_t(\omega)b_t(\omega)) - R_{t+1}^A(\omega)a_t(\omega) \right) \right]. \quad (16)$$

The firm earns returns by investing borrowed funds in the government bond and uses these to pay the returns on equity to households. When taking the outside option, a haircut ξ is applied to the yield R_t^B on the government bond. The term $\Theta_t(\omega)$ reflects that only a fraction $\Theta_t(\omega)$ of the funds borrowed in the inter-firm market can be retained by the borrower when diverting the funds. Notice that this firm will borrow as much as possible, just as long as government bonds minus a haircut ξ pay a positive return. Therefore, similar to the previous situation, the supply conditions dictate the borrowing capacity of the firms in this segment.

We use the zero-profit condition that holds under every state of nature to size the return paid to households by firms in this segment, i.e.

$$R_{t+1}^A(\omega) = \frac{(R_t^B - \xi) (a_t(\omega) + \Theta_t(\omega)b_t(\omega))}{a_t(\omega)}. \quad (17)$$

4.2.3 Firms Choosing to Lend ($\omega \leq \bar{\omega}_t$)

The problem of the firm that lends in the inter-firm market collapses to

$$\max_{l_t(\omega)|\omega} E_t \left[\beta \frac{\lambda_{ct+1}}{\lambda_{ct}} \left(\rho_t \frac{1 - \mu(\bar{\bar{\omega}}_t)}{1 - \mu(\bar{\omega}_t)} l_t(\omega) + \frac{\mu(\bar{\bar{\omega}}_t) - \mu(\bar{\omega}_t)}{1 - \mu(\bar{\omega}_t)} (1 - \theta) l_t(\omega) - R_{t+1}^A(\omega) a_t(\omega) \right) \right], \quad (18)$$

subject to the constraint that $l_t(\omega) \leq a_t(\omega)$.

The term $\frac{1 - \mu(\bar{\bar{\omega}}_t)}{1 - \mu(\bar{\omega}_t)}$ represents the share of loans to firms that choose to produce, yielding the return ρ_t . The term $\frac{\mu(\bar{\bar{\omega}}_t) - \mu(\bar{\omega}_t)}{1 - \mu(\bar{\omega}_t)}$ represents the share of loans to firms that choose to divert the borrowed funds, yielding the average fraction that can be recovered, $1 - \theta$.¹² The firm pays the household the return on assets $R_{t+1}^A(\omega)$.

¹²Proposition 2 establishes that $\bar{\bar{\omega}}_t \geq \bar{\omega}_t$, ensuring that the share of loans to firms that choose to produce and the share of loans to firms that choose to divert the borrowed funds are non-negative and not greater than 1.

Notice that the revenues of the firm are increasing in the amount of loans it supplies. Thus, the budget constraint will hold with equality:

$$l_t(\omega) = a_t(\omega). \quad (19)$$

We use the zero-profit condition that holds under every state of nature to size the return paid to households by firms in this segment, i.e.

$$R_{t+1}^A(\omega) = \rho_t \frac{1 - \mu(\bar{\omega}_t)}{1 - \mu(\bar{\omega}_t)} + \frac{\mu(\bar{\omega}_t) - \mu(\bar{\omega}_t)}{1 - \mu(\bar{\omega}_t)} (1 - \theta). \quad (20)$$

4.3 Capital-Producing Firms

In period t , competitive capital-producing firms buy capital from the goods-producing firms, repair depreciated capital and build new capital. They sell both the new and refurbished capital next period.

Let I_t^g denote aggregate gross investment expenditures. We introduce quadratic adjustment costs measured in units of investment such that the supply of investment goods is given by:

$$I_t^n = Z_{It} \left[1 - \frac{\phi}{2} \left(\frac{I_t^g}{I_{t-1}^g} - 1 \right)^2 \right] I_t^g. \quad (21)$$

The parameter ϕ governs investment adjustment costs for current production relative to past production. The term Z_{it} is a shock process driven by

$$\ln(Z_{It}) = \rho_I \ln(Z_{It-1}) + \varepsilon_{It}, \quad \varepsilon_{It} \sim \mathcal{N}(0, \sigma_I^2). \quad (22)$$

The aggregate capital stock evolves according to:

$$K_t = I_t^n + (1 - \delta)K_{t-1}, \quad (23)$$

where K_t is the amount of capital allocated to the goods-producing firms.

The capital producing firms are owned by households, and solve the problem

$$\max_{I_{t+i}^g} E_t \sum_{i=0}^{\infty} \beta^i \frac{\lambda_{ct+i}}{\lambda_{ct}} \left\{ Q_{t+i} Z_{I_{t+i}} \left[1 - \frac{\phi}{2} \left(\frac{I_{t+i}^g}{I_{t+i-1}^g} - 1 \right)^2 \right] I_{t+i}^g - I_{t+i}^g \right\}. \quad (24)$$

The first-order condition, derived in Section C.2 of the [online appendix](#) (equation (C.5)), determines the equilibrium price of capital Q_t via a standard Tobin's q condition.

Competition will ensure zero profits for the goods-producing firms. Accordingly, the profits rebated to households as aliquot shares will be given by

$$\Pi_t = Q_t Z_{I_t} \left[1 - \frac{\phi}{2} \left(\frac{I_t^g}{I_{t-1}^g} - 1 \right)^2 \right] I_t^g - I_t^g. \quad (25)$$

4.4 The Government

The government finances its transfers T_t by issuing government bonds B_t^G to balance its budget period by period:

$$T_t = B_t^G - R_{t-1}^B B_{t-1}^G. \quad (26)$$

The household and firms can buy government bonds, so

$$B_t^G = B_t^H + D_t, \quad (27)$$

where D_t denotes firms' holdings of government bonds. We assume that the government does not make bonds available to households but sells them only to firms, so that

$$B_t^H = 0 \quad (28)$$

and price the government bond using the households' stochastic discount factor. Although firms hold the bonds in equilibrium, the household's Euler equation for bonds (equation (C.4) in Section C.1 of the [online appendix](#)) still pins down the risk-free rate R_t^B because households are the ultimate owners of all assets in this economy and must be indifferent at the margin.

It implies that the government budget constraint can be written as:

$$T_t = D_t - R_{t-1}^B D_{t-1}. \quad (29)$$

Since firms that take the outside option are subject to a haircut cost ξ on their investment in government bonds in the previous period, the amount of transfers rebated to the household to ensure that there are no deadweight losses in the economy is equal to

$$\Xi_t = \xi D_{t-1}. \quad (30)$$

5 Analytical Characterization of the Equilibrium

This section characterizes the equilibrium analytically and derives its key implications for aggregate output and measured TFP. Section 5.1 provides conditions that pin down the loan rate ρ_t and two cutoff points $\bar{\omega}_t$ and $\bar{\bar{\omega}}_t$. These cutoff points sort firms into three segments depending on the realization of the idiosyncratic productivity shock: 1) firms with $\omega \leq \bar{\omega}_t$ lend; 2) firms with $\bar{\omega}_t < \omega < \bar{\bar{\omega}}_t$ borrow and default; 3) firms with $\omega \geq \bar{\bar{\omega}}_t$ borrow and produce. Our strategy is to posit a solution that satisfies the first-order conditions of the intermediate goods sector and then verify that no firm can or has an incentive to switch to a different segment. To this purpose, we formulate three propositions. The aggregation of individual firm decisions and the formal definition of competitive equilibrium are left for Section C.3 of the [online appendix](#). Section 5.2 then integrates individual firm outputs across the producing segment to obtain aggregate output, and shows how the equilibrium cutoff $\bar{\bar{\omega}}_t$ generates an endogenous component of the measured Solow residual.

5.1 The Loan Rate and Cutoff Points

Here we clarify the importance of the screening technology that allows lenders to tell whether borrowers can be expected to make more than the outside option. This technology prevents the firms with the lowest firm-specific technology levels from borrowing.

Assume (and then verify) that if the firms that have a firm-specific productivity $\omega < \bar{\omega}_t$

will choose to lend, then the following condition holds:

$$E_t \left[\beta \frac{\lambda_{ct+1}}{\lambda_{ct}} \left(\frac{\kappa_{t+1}}{Q_t} \bar{\omega}_t + \frac{(1-\delta)}{Q_t} Q_{t+1} \right) \right] = E_t \left[\beta \frac{\lambda_{ct+1}}{\lambda_{ct}} (R_t^B - \xi) \right]. \quad (31)$$

This condition imposes that for a firm with productivity right at the cutoff point between borrowing and lending, $\omega = \bar{\omega}_t$, the expected return from producing $R_{t+1}(\bar{\omega}_t)$ equals the expected return of the outside option $R_t^B - \xi$.

Additionally, the following condition applies to firms that lend. We need to compare the expected return from lending to the expected return from diverting funds. For a firm whose productivity lies exactly at the cutoff between borrowing and lending, $\omega = \bar{\omega}_t$, the expected return from lending, i.e., the lending rate net of the losses expected from firms diverting funds, matches the expected return from borrowing and defaulting:

$$\begin{aligned} E_t \left[\beta \frac{\lambda_{ct+1}}{\lambda_{ct}} \left(\rho_t \frac{1 - \mu(\bar{\omega}_t)}{1 - \mu(\bar{\omega}_t)} + (1 - \theta) \frac{\mu(\bar{\omega}_t) - \mu(\bar{\omega}_t)}{1 - \mu(\bar{\omega}_t)} \right) a_t \right] = \\ E_t \left[\beta \frac{\lambda_{ct+1}}{\lambda_{ct}} (R_t^B - \xi) (a_t + \Theta_t(\bar{\omega}_t) b_t) \right]. \end{aligned} \quad (32)$$

Notice there is no dependence of b_t on ω because $b_t = b_t(\omega)$ for all ω due to asymmetric information. At this point, we have spelled out all the elements that are needed for our first proposition.

Proposition 1. *Given that $\Theta_t(\omega)$ is non-negative and increasing in ω , firms with idiosyncratic productivity ω less than the cutoff point $\bar{\omega}_t$ will have no incentive to deviate from lending.*

A proof of this proposition is offered in Section A of the [online appendix](#). The proof relies primarily on equations (31) and (32).

For the next step of pinning down $\bar{\omega}_t$, consider the following condition:

$$E_t \left[\beta \frac{\lambda_{ct+1}}{\lambda_{ct}} (R_t^B - \xi) (a_t + \Theta_t(\bar{\omega}_t^*) b_t) \right] = E_t \left[\beta \frac{\lambda_{ct+1}}{\lambda_{ct}} (R_{t+1}(\bar{\omega}_t^*) (a_t + b_t) - \rho_t b_t) \right]. \quad (33)$$

It imposes that the marginal firm with productivity level $\bar{\omega}_t^*$ will be indifferent between diverting funds and producing.

Taking the first derivative with respect to $\bar{\omega}_t^*$ of both sides of equation (33), we will also need to impose the following condition on these derivatives:

$$E_t \left[\beta \frac{\lambda_{ct+1}}{\lambda_{ct}} (R_t^B - \xi) \Theta'_t(\bar{\omega}_t^*) b_t \right] < E_t \left[\beta \frac{\lambda_{ct+1}}{\lambda_{ct}} (R'_{t+1}(\bar{\omega}_t^*) (a_t + b_t)) \right]. \quad (34)$$

In other words, for a marginal firm with the productivity level $\bar{\omega}_t^*$ the expected returns from diverting funds grow less steeply than the expected returns from producing. This condition is really an additional condition on the choice of the function $\Theta_t(\omega)$. It implies that a firm with a slightly higher productivity than $\bar{\omega}_t^*$ will prefer producing to diverting funds, while a firm with a slightly lower productivity than $\bar{\omega}_t^*$ will prefer diverting funds to producing. Note that it is a local condition in the sense that it applies to a marginal firm with the productivity level $\bar{\omega}_t^*$.

We define $\bar{\bar{\omega}}_t = \max(\bar{\omega}_t, \bar{\omega}_t^*)$. This definition ensures that terms that represent shares in the expression of the return from lending, for instance, $\frac{\mu(\bar{\omega}_t) - \mu(\bar{\omega}_t^*)}{1 - \mu(\bar{\omega}_t)}$ in equation (32), are well-defined. But we can do more and establish that $\bar{\bar{\omega}}_t$ is never less than $\bar{\omega}_t$, in line with the following proposition.

Proposition 2. *Take $\bar{\omega}_t^*$ to be the productivity of a firm indifferent between diverting funds and producing. Given that $\Theta_t(\omega)$ is non-negative, then it must be that $\bar{\bar{\omega}}_t \geq \bar{\omega}_t$.*

Notice that, by definition of $\bar{\bar{\omega}}_t$, a corollary of Proposition 2 is that $\bar{\bar{\omega}}_t = \bar{\omega}_t^*$, so $\bar{\bar{\omega}}_t \geq \bar{\omega}_t$.

Before moving to our final proposition, we also need to choose a function such that

$$E_t \left[\beta \frac{\lambda_{ct+1}}{\lambda_{ct}} (R_t^B - \xi) \Theta'_t(\bar{\omega}_t) b_t \right] < E_t \left[\beta \frac{\lambda_{ct+1}}{\lambda_{ct}} \frac{\kappa_{t+1}}{Q_t} (a_t + b_t) \right]. \quad (35)$$

The left-hand side is the maximum expected marginal gain from diversion, evaluated at the marginal firm $\omega = \bar{\omega}_t$ that is indifferent between diverting funds and lending. Because of the concavity of the function, $\Theta_t(\bar{\omega}_t)$ is at least as large as for any other firm with $\omega > \bar{\omega}_t$. The right-hand side can also be written as $E_t \left[\beta \frac{\lambda_{ct+1}}{\lambda_{ct}} R'_{t+1}(\omega) (a_t + b_t) \right]$, the expected marginal gain from producing; since $R_{t+1}(\omega) = (\kappa_{t+1}/Q_t) \omega + (1 - \delta) Q_{t+1}/Q_t$, its derivative at ω equals κ_{t+1}/Q_t which does not depend on ω . Because the expected marginal benefit from diversion is decreasing in ω (as Θ_t is concave), if this global condition holds at $\omega = \bar{\omega}_t$, then it implies

the local condition (34). We can use this global condition to prove that no diverting firm has an incentive to deviate by producing, and no producing firm has an incentive to switch to diverting. With this condition specified, we can state our final proposition.

Proposition 3. *Given that $\Theta_t(\omega)$ is non-negative, concave and increasing in ω , equations (31), (32), and (33), together with the slope condition under (35) are sufficient to ensure that, depending on the realization of their idiosyncratic productivity ω , firms sort themselves into three groups:*

1. *firms with $\omega \leq \bar{\omega}_t$ lend;*
2. *firms with $\bar{\omega}_t < \omega < \bar{\bar{\omega}}_t$ borrow and default;*
3. *firms with $\omega \geq \bar{\bar{\omega}}_t$ borrow and produce.*

A proof of this proposition is in Section A of the [online appendix](#).

5.2 Aggregate Output and the TFP Decomposition

Once the propositions above pin down the equilibrium cutoffs, individual producing-firm outputs can be integrated across the mass of producing firms. The full derivation is in Section C of the [online appendix](#). The key step combines two facts: (i) all producing firms hold the same capital \tilde{k}_t (credit rationing); and (ii) because the labor commitment is also made before ω is observed, each firm's optimal labor satisfies $h_{t+1}(\omega) = \tilde{h}_{t+1}$ from (11). Individual output is therefore $y_{t+1}(\omega) = Z_{t+1} \omega \tilde{k}_t^\alpha \tilde{h}_{t+1}^{1-\alpha}$, which is linear in ω . Integrating across the producing segment yields:

$$Y_t = Z_t K_{t-1}^\alpha H_t^{1-\alpha} \frac{\int_0^1 \omega \mu'(\omega) d\omega}{1 - \mu(\bar{\bar{\omega}}_{t-1})}. \quad (36)$$

This equation reveals the central insight: measured TFP—the Solow residual defined by $Y_t/(K_{t-1}^\alpha H_t^{1-\alpha})$ —equals the exogenous shock Z_t multiplied by $E[\omega \mid \omega \geq \bar{\bar{\omega}}_{t-1}]$, the conditional mean of idiosyncratic productivity among producing firms. Because the cutoff $\bar{\bar{\omega}}_{t-1}$ responds to financial conditions, the Solow residual contains an endogenous component even

when Z_t is held fixed: a tightening that raises $\bar{\omega}_{t-1}$ selects more productive firms into operation, lifting $E[\omega \mid \omega \geq \bar{\omega}_{t-1}]$. Evaluating the integral in closed form requires a distributional assumption on μ .

5.2.1 Functional Form for the Distribution of Idiosyncratic Productivity

We choose the Beta distribution to govern draws of the idiosyncratic productivity $\omega \sim \text{Beta}(\eta_1, \eta_2)$ on $[0, 1]$. With shape parameters η_1 and η_2 , the probability density function of ω is

$$\mu'_{\eta_1, \eta_2}(\omega) = \frac{\omega^{\eta_1-1}(1-\omega)^{\eta_2-1}}{B(\eta_1, \eta_2)}, \quad (37)$$

where $B(\eta_1, \eta_2) = \Gamma(\eta_1)\Gamma(\eta_2)/\Gamma(\eta_1 + \eta_2)$ and Γ is the Gamma function.

We also need to specify the diversion function $\Theta_t(\omega)$, which governs the fraction of borrowed funds $b_t(\omega)$ that a firm can divert. To align our functional choice with the structural justification for this function derived in Section B of the [online appendix](#), this function must be non-negative, increasing, and concave in ω . Among the many functional forms consistent with these requirements, we choose the power specification

$$\Theta_t(\omega) = \omega^\psi F_t, \quad (38)$$

where ψ is a parameter and F_t is determined endogenously to satisfy

$$\theta = \int_{\bar{\omega}_t}^{\bar{\bar{\omega}}_t} \omega^\psi F_t \frac{\mu'(\omega)}{\mu(\bar{\bar{\omega}}_t) - \mu(\bar{\omega}_t)} d\omega. \quad (39)$$

Section C.9 of the [online appendix](#) provides the functional form of F_t .¹³ We have verified numerically that the choices of ψ and other parameters that influence F_t satisfy the slope condition (35), which is used in the proof that ω determines each firm's segment.¹⁴

¹³Section C.9 of the [online appendix](#) also shows that $F_t = \frac{\theta}{\bar{\omega}_t^\psi}$ in the limiting case as $\bar{\bar{\omega}}_t \rightarrow \bar{\omega}_t$.

¹⁴The equation for F_t makes $\Theta(\omega)$ depend on the equilibrium cutoffs, so a firm with a given ω can have a different diversion share when the defaulting interval changes. Individual firms do not internalize this general-equilibrium effect. For the model developed in Section B of the [online appendix](#), we show that a general equilibrium mechanism compresses the diversion rate when more firms attempt to divert resources. Numerically, we verify that the chosen diversion function shares the same property in the steady state of the

5.2.2 Closed-form Decomposition of TFP into Endogenous and Exogenous Components

With the Beta distribution, the integral in equation (36) evaluates in closed form. The key step rewrites $\omega \cdot \mu'_{\eta_1, \eta_2}(\omega) \propto \omega^{\eta_1} (1 - \omega)^{\eta_2 - 1}$ as a scalar multiple of the $\text{Beta}(\eta_1 + 1, \eta_2)$ density kernel, converting the integral via the ratio of Beta functions $B(\eta_1 + 1, \eta_2)/B(\eta_1, \eta_2) = \eta_1/(\eta_1 + \eta_2)$. Define

$$c_\alpha \equiv \frac{B(\eta_1 + 1, \eta_2)}{B(\eta_1, \eta_2)} = \frac{\eta_1}{\eta_1 + \eta_2}, \quad (40)$$

which is the unconditional mean of $\omega \sim \text{Beta}(\eta_1, \eta_2)$. The full derivation is provided in Section C of the [online appendix](#), equation (C.53). Aggregate output then simplifies to

$$Y_t = \Phi_t \cdot Z_t \cdot K_{t-1}^\alpha H_t^{1-\alpha}, \quad (41)$$

where the endogenous TFP prefactor is

$$\Phi_t \equiv c_\alpha \frac{1 - \mu_{\eta_1+1, \eta_2}(\bar{\omega}_{t-1})}{1 - \mu_{\eta_1, \eta_2}(\bar{\omega}_{t-1})}. \quad (42)$$

The right-hand side equals $E[\omega \mid \omega \geq \bar{\omega}_{t-1}]$, the conditional expectation of ω among producing firms. We retain the notation c_α for consistency with the appendix; here it collapses to the unconditional Beta mean $\eta_1/(\eta_1 + \eta_2)$.

Equations (41)–(42) give the TFP decomposition a precise quantitative form. The measured Solow residual is

$$\text{TFP}_t \equiv \frac{Y_t}{K_{t-1}^\alpha H_t^{1-\alpha}} = \underbrace{\Phi_t}_{\text{endogenous}} \cdot \underbrace{Z_t}_{\text{exogenous}}.$$

As the economy approaches the frictionless RBC limit, $\bar{\omega}_t \rightarrow 1$ for all t ; consequently $\Phi_t \rightarrow 1$ (see equation F.28 in Section F of the [online appendix](#)), and the Solow residual moves one-for-one with Z_t . In our model, $\bar{\omega}_{t-1}$ responds endogenously to financial conditions, causing Φ_t to fluctuate. A tightening that raises the default cutoff concentrates production among higher-productivity firms and raises Φ_t ; a loosening has the opposite effect. The Solow residual

model.

therefore inherits variation from both the exogenous technology process and the endogenous reallocation of production across firms. Section 7.1 and the quantitative analysis document the empirical importance of this channel.

6 Parameter Choices and Model Solution

We now turn to the quantitative implementation of the model. This section is organized as follows. Section 6.1 describes the choice of parameter values, combining steady-state calibration targets with a simulated method of moments estimation of the shock processes and adjustment-cost parameter. Section 6.2 outlines the second-order perturbation solution method. Section 6.3 reports the targeted moments to assess the overall fit of the model. We leave a discussion of untargeted moments for Section 7.3.

6.1 Parameters

We choose model parameters with a mix of estimation and calibration. Specifically, we estimate model-specific parameters using the simulated method of moments (SMM) on second moments. We pin down other parameters by matching first moments. Finally, we set a few parameters to standard values from the literature.

We use the simulated method of moments (SMM) to size the shock processes, total factor productivity, Z_t , investment-specific technology, Z_{It} , and the consumption shock ν_{ct} . We model each shock as an AR(1) process. We size the persistence parameters and standard deviations of the innovations. Along with sizing these shocks, we also estimate by SMM the investment adjustment cost parameter, ϕ .

For the implementation of the SMM estimation, we minimize a quadratic objective function based on variances, correlations, and autocorrelations of real GDP, real consumption, the relative price of investment, and the average delinquency rate on business loans at commercial banks.¹⁵

¹⁵For GDP and relative prices we used data from the NIPA Release of the U.S. Bureau of Economic Analysis. We use chain-type indexes. The relative price of investment is the ratio of the price index for gross private domestic investment to the price index for personal consumption excluding food and energy. The delinquency rate on business loans at commercial banks is from the data release Charge-Off and Delinquency

Our data run from the first quarter of 1987 through the fourth quarter of 2021, matching the sample period in Section 3, a choice dictated by the availability of data on the dispersion of productivity. The data moments for the SMM exercise are computed after HP-filtering the data (using a value for the smoothing parameter of 1,600 as standard for quarterly data). We take the same approach for observed and simulated data. For the SMM objective function, we employ a modified optimal weighting matrix with model moments from the average of 10 simulated samples of the same length as the observed data.¹⁶

We choose selected parameters to match steady-state targets for key variables, using first moments from the same sample period as the SMM estimation. We size the two parameters of the Beta distribution η_1 and η_2 , the haircut ξ , and the diversion share θ at values that jointly match the 1) within-industry dispersion of productivity as captured by the weighted average gap in TFP between the 90th and 10th percentile firm, 2) the average spread between the return on assets for households and the risk-free rate of the outside option of firms, 3) the average share of bank credit, and 4) the average delinquency rate on business loans.¹⁷ The idea behind these choices of the parameters is the following. The shape of the productivity distribution is important to match that the establishment at the 90th percentile is 2.08 as productive as the one at the 10th percentile. We compute the target ratio of weighted-average within-industry dispersion over the sample period by taking the mean of the data points underlying the solid line in the top panel of Figure 1. Section E of the [online appendix](#) provides formulas to calculate this target ratio from the model. We map lending firms into stylized banks, which provide about 47% of private credit in the United States.¹⁸ We match

Rates on Loans and Leases at Commercial Banks of the Board of Governors of the Federal Reserve System. The series for real GDP and delinquency rates are the same as those used in Section 3. In our simple model, there is no distinction between delinquency rates, charge-off rates, and the outright default rate of borrowers.

¹⁶The standard optimal weighting matrix assigns a weight proportional to the inverse of the long-run variance of each sample moment. The variance of consumption is estimated very tightly from the data; so, the SMM optimal weighting matrix assigns a disproportionate weight to that moment to the detriment of pinning down other important moments —the variance of the investment price, the variance of the default rate, and the correlation between GDP and the default rate. To prevent this single-moment dominance, we increase the weight on the GDP–default-rate correlation to the full weight of the consumption variance, while while the other two moments receive half the weight of the consumption variance.

¹⁷The shape of the productivity distribution affects all the described targets, so the word “*jointly*” follows.

¹⁸We construct our measure using data from the Z.1 Financial Accounts of the United States. We sum non-financial corporate and non-financial non-corporate business loans and subtract non-financial business other loans and advances to form the numerator. We divide by total credit to the non-financial private sector from the Bank for International Settlements and take the average. We note that the BIS denominator covers a broader set of borrowers than the non-financial non-corporate sector included in the Z.1 numerator, so our

this share to our model counterpart expression $\frac{L_t}{B_t^{\text{tot}}}$ evaluated in the steady state. Because all external funds in the model are intermediated through lending firms, $\frac{L_t}{B_t^{\text{tot}}}$ is the model analogue of the share of firm financing provided through the banking system. We use the same data on delinquent loans as in Section 3 to size the average delinquency rate at 2.6%. We map it to χ_t , the share of borrowing firms diverting funds,¹⁹

$$\chi_t \equiv \frac{\mu_{\eta_1, \eta_2}(\bar{\omega}_t) - \mu_{\eta_1, \eta_2}(\bar{\omega}_t)}{1 - \mu_{\eta_1, \eta_2}(\bar{\omega}_t)}, \quad (43)$$

and calibrate θ so that its steady-state value matches the 2.6% target. We match the average spread between the bank prime loan rate and a one-year Treasury bill to the haircut parameter ξ . At the deterministic steady state, $R^A = R^B$ from the household Euler equations, so the spread between the return to assets for households and the outside-option return $R^B - \xi$ equals exactly ξ ; equivalently, ξ is the model counterpart of the observed prime-rate–Treasury-bill spread. The prime rate is a good match for R_t^A , since it is an unsecured (defaultable) rate.²⁰

Before discussing SMM-estimated parameters, we first outline the calibration of the remaining parameters. Calibrated values appear at the top of Table 2 and follow guidance from prior literature: the discount rate $\beta = 0.9925$, the capital share $\alpha = 0.3$, the depreciation rate $\delta = 0.01$, the inverse Frisch elasticity of labor supply $\nu = 2$. Under KPR preferences, the utility function features log utility in consumption, so $\sigma = 1$ is imposed implicitly and is not a free parameter.

We set ϑ at 0.84, a value that ensures that aggregate labor takes a value of 1 in the steady state.

measure should be interpreted as an approximation.

¹⁹Our model does not distinguish between delinquency, charge-offs, and outright default: a diverting firm fails to repay in full, which we treat as a single default event. We therefore take the observed delinquency rate on business loans as a proxy for the model’s default rate χ_t .

²⁰The data for the bank prime loan rate and the return on one-year Treasury bills are from the Federal Reserve Board’s H.15 Release.

6.2 Model Solution

We solve the model with an exact second-order perturbation method using the pruning algorithm of Kim et al. (2008). When drawing data samples from the model, we found that a first-order perturbation method was not sufficiently accurate to avoid the upper cutoff point $\bar{\omega}_t$ falling below the lower cutoff point $\underline{\omega}_t$, a theoretical impossibility. By contrast, a second-order perturbation solution is sufficiently accurate to avoid this problem. We provide further accuracy checks in Section H of the [online appendix](#).

6.3 Simulated Method of Moments, Results

Table 3 reports confidence intervals for both the data moments and the model moments alongside the point estimates. The model matches most of the targeted second moments well: the data and model confidence intervals overlap for 10 of the 14 moments. The four variances are all matched within sampling uncertainty, as are the correlations involving GDP, consumption, and the default rate, and the autocorrelations of GDP and consumption. For four moments—the correlation between GDP and the investment price, the correlation between the investment price and the default rate, and the first-order autocorrelations of the investment price and the default rate—the confidence intervals do not overlap, indicating a quantitatively significant discrepancy. Nonetheless, for every moment in the table, including these four, the model correctly matches the sign.

Table 4 reports the relative contributions of the shocks in our calibration to fluctuations in variables targeted by the SMM procedure and other key variables. Aggregate TFP shocks are the primary drivers of fluctuations in GDP, but we find that consumption shocks are the most important driver of TFP dispersion. The table also shows that IST shocks are less important for our model to hit the moments targeted in the SMM estimation procedure.

7 Dynamic Responses to Shocks and Model Assessment

Before turning to an assessment of the model by checking untargeted moments, Section 7.1 discusses the effects of an aggregate TFP shock, and Section 7.2 considers consumption and investment specific shocks. We compare our model to a frictionless RBC model, a special case of our model in which the most efficient firm is the only one to produce and attracts all household savings.²¹ Whereas the first two parts of this section focus on moments conditional on specific shocks, the last part turns to unconditional moments. We compare moments untargeted by our calibration using observed and model synthetic data.

7.1 Effects of a TFP Shock

Figure 2 shows the responses of key aggregate variables to a one-standard-deviation expansionary TFP shock, ε_t^z . The size of the shock process and its persistence follow the estimates from Section 6. In each panel, the solid lines show responses from the baseline model with financial frictions and endogenous TFP dispersion. The dashed lines show responses in a frictionless RBC model, a special case of our model in which all production is carried out by the most efficient firm.

On impact (period 1 in the figure), the positive productivity shock increases the marginal product of capital, which in turn raises the average return on capital. It is this standard exogenous productivity that pushes up aggregate output on impact. An endogenous TFP channel kicks in from the second period.

Since the shock is autocorrelated but decaying ($\rho_z < 1$), the expected marginal product of capital for the next period is lower relative to the impact response. In line with this effect, the demand for capital after initially expanding starts declining. This expected decline brings down the expected price of capital. Accordingly, the average return to capital falls. This return is averaged across firms with different idiosyncratic productivity, and it includes a term that reflects the marginal product of capital and another term that captures gains or

²¹Section F of the [online appendix](#) proves that when firms can credibly disclose their productivity and $\xi = 0$, our model collapses to a standard RBC model.

losses from the resale of undepreciated capital. It is this second term that pushes down the average return on capital.

As a reminder, Equation (33) pins down $\bar{\omega}_t$.²² From this equation, we can see that, all else equal, a decrease in the average return of capital pushes up $\bar{\omega}_t$. While many of the variables entering equation (33) do change, the decrease in the average return of capital is the most important quantitatively. Intuitively, more productive firms are better able to withstand a decrease in the average return to capital. As $\bar{\omega}_t$ rises, the productivity dispersion declines, as captured by a reduction in the gap between the productivity of the 90th and 10th percentile firms.

As the aggregate technology shock is expected to gradually wane, the expected return to producing falls for the marginal firm—chiefly because the expected resale price of capital declines, the same force that lowers the average return on capital and raises $\bar{\omega}_t$. The risk-free outside-option return R_t^B declines as well, but by less, so the firm indifferent between lending and borrowing tips toward lending and the lower cutoff $\bar{\omega}_t$ rises.

Notice that as $\bar{\omega}_t$ does not rise quite as much as $\bar{\omega}_t$, the default rate χ_t defined in equation (43) goes down. Accordingly, default rates move countercyclically. More generally, expansionary shocks can lead to an increase or a decrease in the default rate depending on whether the effects on the expected returns dominate or the effects through the outside option dominate.

It is useful to keep the structural roles of the two cutoffs distinct, in anticipation of the comparison developed in Section 8. The upper cutoff $\bar{\omega}_{t-1}$ enters aggregate production directly: it is the threshold above which firms choose to produce and repay; firms with $\bar{\omega}_{t-1} < \omega < \bar{\omega}_{t-1}$ borrow and default, so only firms with $\omega \geq \bar{\omega}_{t-1}$ contribute to output in period t . It is therefore $\bar{\omega}_{t-1}$ that governs the endogenous prefactor Φ_t and shapes measured TFP. The lower cutoff $\bar{\omega}_t$, by contrast, is purely forward-looking: it marks the boundary between lenders and all borrowers—producers and defaulters alike are on the borrowing side of this threshold—and responds to current and expected-future conditions, but it does not appear directly in the aggregate production function. The two cutoffs co-move in equilibrium, yet their separate determination allows the production-relevant margin ($\bar{\omega}$) and the lending-

²²The aggregate counterpart of equation (33) is given in Section C of the [online appendix](#).

incentive margin ($\bar{\omega}$) to absorb shocks with some degree of independence. This independence is precisely what is lost in the no-default model, where a single cutoff must play both roles simultaneously.

Due to a fall in the productivity dispersion, investment rises more than in a frictionless RBC model. Total output also increases by more. On impact, consumption rises relatively less compared to a frictionless RBC model as consumption is crowded out by more productive investments. Subsequently, starting from the next period, consumption rises relatively more due to higher total output.

Summing up, financial frictions and firm dispersion lead to persistent output and interest rate dynamics.

7.2 Effects of Other Shocks

Figure 3 shows the responses of key aggregate variables to a one-standard-deviation shock to consumption preferences, $\varepsilon_{\nu t}$, leading to an expansion in consumption. Again, the size of the shock process and its persistence follow the estimates from Section 6.

The differences between our model and the canonical RBC model are more dramatic for this shock. Starting from the risk-free rate, its rise is supercharged relative to the rise in the RBC model. Two forces cumulate to produce this difference, a substitution and a wealth effect. By construction, the shock has a pronounced substitution effect, which results in a rise in the risk-free rate in both models. In a closed-economy setting which precludes the possibility of borrowing from abroad, the rise is necessary to push down investment leaving room to consume more today for any level of labor input. But in our model with financial frictions, this rise in the risk free rate also affects the outside option for defaulting firms. There is a rise in borrowing costs and an increase in defaults. The rise in borrowing costs also contributes to concentrating production in higher productivity firms—the ones more able to withstand increases in borrowing costs. As a result, the productivity dispersion falls. This fall in dispersion is the extra endogenous kick to productivity which is the main determinant of the difference between the evolution of output across the two models. But consider that this effect generates a wealth effect that boosts consumption by more, hence the need for a much larger rise in the risk-free interest rate in our model relative to the frictionless model.

Although the differences between the two models are quite persistent, we have confirmed that both models are stationary. The model with financial frictions does return to its steady state at horizons well beyond the one for the figure.

Figure 4 shows the responses of key aggregate variables to an expansionary one-standard-deviation shock to investment technology ε_{It} . And once again, the size of the shock process and its persistence follow the estimates from Section 6.

Investment-specific technology shocks generate strong substitution effects as the benefits of higher productivity can only be accrued by boosting investment. Accordingly very similar considerations apply as for consumption shocks. The real interest rate needs to rise to curb consumption in this case. Default rates rise, as the rise in borrowing costs is the dominant force, dispersion decreases, and total factor productivity rises endogenously, giving rise to a sizable wealth effect (relative to the size of the shock). As Section 8 makes clear, the strategic-default margin plays a dual role throughout: it acts as an amplifier by concentrating production among high-productivity firms, and simultaneously as a stabilizer that prevents the oscillatory TFP dynamics that emerge when a single cutoff must govern both the selection of producers and the lending incentive.

7.3 Untargeted Moments

Having traced the model’s response to each shock, we now turn to a set of regression-based untargeted moments that ask whether the same conditional correlations apparent in the U.S. data also arise in the model.

Table 5 reports estimates from four ordinary least squares regressions on annual U.S. data from 1987 to 2021. In columns 1 and 2, the dependent variable is the weighted-average within-industry 90–10 TFP gap from Figure 1, expressed in percentage points.²³ Column 1 regresses that gap on the contemporaneous growth rate of real GDP per capita (in annual percent) and the delinquency rate on business loans at commercial banks (in annual percent); column 2 adds one lag of the growth rate.²⁴ In columns 3 and 4, the dependent variable is the

²³This is the value-added-weighted 90–10 measure plotted by the solid line in the top panel of Figure 1, built from the BLS Dispersion Statistics on Productivity (Cunningham et al., 2023) together with industry value-added weights from the Bureau of Economic Analysis.

²⁴Real GDP per capita is the ratio of real gross domestic product (in chained 2017 dollars) to mid-

delinquency rate itself, regressed on its own one-period lag, contemporaneous GDP per capita growth, and—in column 4—an ex-ante one-year real interest rate (in annual percent).²⁵ We are deliberately humble about what these regressions deliver: they are partial correlations, not causal estimates, and we use them as additional untargeted moments against which to judge the calibration.

The point estimates in columns 1 and 2 of Table 5 confirm the pattern visible in Figure 1. A one-percentage-point higher annual growth rate is associated with a 2.86 pp narrower 90–10 gap in column 1, and the lagged-growth coefficient in column 2 is also negative and significantly different from zero. Both columns also show a sizable negative coefficient on the delinquency rate (about -4.2 pp). For the results shown in columns 3 and 4, consider that the delinquency rate is persistent (lagged coefficient around 0.8), it falls when growth is high (by about 0.22 pp per percentage point of growth), and—once the real rate is added in column 4—it rises with the real rate at a coefficient of 0.14 pp. Read together, these last two columns simply summarise the empirical fact that delinquencies are higher when growth is depressed and when real rates are high.

To assess the model along these dimensions, we run exactly the same four regressions on simulated data. For each of 1,000 replications, we simulate the model at quarterly frequency over a span that matches the length of the observed sample and then aggregate to annual frequency by averaging within each year, mirroring the construction in the data. The model’s analogue of the dispersion measure is the time series of $100 \cdot \omega_{90}/\omega_{10}$ built from the upper cutoff $\bar{\omega}_t$; the analogue of the delinquency rate is the model default rate times one hundred; growth is the annual log-difference of yearly model GDP per capita, expressed in percent; and the analogue of the real rate is the net annualised return on bonds, $400 \cdot (R^B - 1)$, also in percent. Lags are constructed exactly as in the data and we estimate the four specifications

year population, both from the U.S. Bureau of Economic Analysis Gross Domestic Product release within the National Income and Product Accounts; the population concept covers resident population plus armed forces overseas. The delinquency rate on business loans at commercial banks comes from the Federal Reserve Board’s Charge-Off and Delinquency Rates on Loans and Leases at Commercial Banks release. We aggregate the underlying quarterly delinquency series to annual frequency by averaging within the year.

²⁵The ex-ante one-year real interest rate is from the Federal Reserve Bank of Cleveland’s Inflation Expectations release; the underlying model uses Treasury yields, inflation data, inflation swaps, and survey-based measures of inflation expectations. The original monthly series is aggregated to annual frequency by averaging within the year.

by OLS. Table 6 reports the average slope coefficient across the 1,000 replications, with the 5th and 95th percentiles in brackets.

In the dispersion regressions (columns 1 and 2 of Table 6), the model produces a negative and significant coefficient on current GDP per capita growth and on its lag, matching the sign of the data estimates in Table 5. The contemporaneous delinquency coefficient is also negative in the model, and its point estimate (about -4.2) is essentially the same as in the data, although the model's 90% range is wide enough to include zero. The dynamics behind these signs are visible in the impulse responses of Section 7.1: in Figure 2, the same TFP shock that raises output narrows the 90–10 gap and lowers the default rate, so dispersion, growth, and delinquencies co-move in a way that produces negative partial correlations period by period. The wide range on the delinquency coefficient in columns 1 and 2 is a signal that the model captures the sign of this conditional correlation more sharply than its magnitude across small simulated samples.

The delinquency regressions (columns 3 and 4 of Table 6) provide a sanity check on a key contribution of the paper. The lagged delinquency coefficient is positive and significant, as in the data; the contemporaneous growth coefficient in column 3 is negative and significant (-0.11 in the model versus -0.22 in the data); and the real-interest-rate coefficient in column 4 is positive and tightly estimated (0.26 in the model versus 0.14 in the data). These signs line up with the dynamics shown in Sections 7.1 and 7.2: the consumption and investment-specific shocks of Figures 3 and 4 push the real rate up and the default rate up at the same time; the TFP shock of Figure 2 raises growth and pulls delinquencies down. The one cell that does not line up is the growth coefficient in column 4, which is essentially zero in the model: once the real rate enters the regression, it absorbs the cyclical variation in delinquencies that growth carries on its own in column 3.

Taken together, these untargeted regressions reproduce, in sign, four of the partial correlations that we read off the data, including the central one that delinquencies rise when growth is depressed and when real rates are high. Section 8 now asks how much of the endogenous TFP behaviour underlying these regressions depends specifically on the strategic-default margin.

8 Endogenous TFP Dynamics With and Without Strategic Default

To sharpen the intuition for how financial frictions shape measured TFP, we contrast the baseline model with a simplified variant in which strategic default is ruled out by assumption. The comparison between these two model variants allows us to showcase the important role of default in driving the endogenous component of TFP.

In the *no-default model*, the outside option is unavailable, so no firm has an incentive to divert borrowed funds; the two-cutoff structure of the baseline collapses to a single cutoff $\bar{\omega}_t$ that separates lenders from producers. All other model blocks—the household problem, capital-producing firms, and the shock processes—are unchanged. Because no funds are diverted, the government sector drops out entirely. Section G of the [online appendix](#) develops this variant in full detail.

Figure 5 presents a side-by-side comparison of the TFP decomposition for the two models across all three shocks keeping the model parameters unchanged across models. In each panel, the stacked areas display the two components of the Solow residual introduced in Section 5.2: the blue area is the exogenous component $\log Z_t$, the orange area is the endogenous prefactor $\log \Phi_t$, and the top of the stack traces total measured \log TFP. The left column shows the baseline (two-cutoff) model; the right column shows the no-default (one-cutoff) model. The y -axis scale is shared within each row so that magnitudes are directly comparable across the two models.

For the TFP shock (top row), both models generate a positive TFP response, with the endogenous component amplifying the exogenous shock. The amplification is stronger in the baseline model: strategic default provides an additional margin through which the upper cutoff $\bar{\omega}_t$ responds to aggregate conditions, concentrating production among high-productivity firms and lifting the TFP prefactor further than the single-cutoff model can. For the preference and IST shocks (middle and bottom rows), Z_t is exactly zero—these shocks do not move the exogenous productivity process—so the entire TFP response is endogenous and confined to Φ_t . The contrast between the two models is striking here: the baseline model generates a smooth, monotonically decaying Φ_t response, while the no-default model

generates a pronounced oscillatory pattern that alternates sign period by period.

8.1 Results for the No-Default Model

Two features of the no-default model interact to produce oscillatory behavior. The first feature is the role of the cutoff as a lagged state variable in production. As shown in equation (G.20) of Appendix G, aggregate output in period t depends on the cutoff from the *previous* period:

$$Y_t = E[\omega \mid \omega \geq \bar{\omega}_{t-1}] \cdot Z_t K_{t-1}^\alpha H_t^{1-\alpha}. \quad (44)$$

The mechanism is straightforward. Because capital is allocated equally across all producing firms, the aggregate capital–labor ratio is identical for every producer. The only source of heterogeneity in individual output is idiosyncratic productivity ω ; summing across producers, aggregate output is therefore proportional to the conditional mean of ω among firms with $\omega \geq \bar{\omega}_{t-1}$. This is the *selection effect*: when the cutoff rises, the producing cohort is leaner and average productivity rises; when the cutoff falls, lower-productivity firms enter the cohort and drag down average productivity.

The second feature is the forward-looking determination of $\bar{\omega}_t$. The indifference condition (G.9) equates the expected discounted return from lending today with the expected discounted return from producing, so the cutoff is a function of current and expected future variables, not of past conditions. This creates a fundamental tension: $\bar{\omega}_t$ responds on impact to current shocks, but its effect on production materialises only one period later, through equation (44).

To fix ideas, consider a positive TFP shock; analogous dynamics set in for the preference and IST shocks, where the initiating impulse comes from the demand side rather than from a direct shift in Z_t . The feedback loop can be traced period by period.

- *Period t .* The positive TFP shock is persistent, so it raises the *expected* marginal product of capital in period $t + 1$. Because the cutoff $\bar{\omega}_t$ is forward-looking—it equates the expected discounted return from producing in $t + 1$ with the outside-option return—this higher expected return makes production more attractive. The indifference condition therefore shifts so that $\bar{\omega}_t$ *falls*, admitting additional lower-productivity producers into

the active cohort.

- *Period $t + 1$.* The lower $\bar{\omega}_t$ reduces $E[\omega \mid \omega \geq \bar{\omega}_t]$, so effective TFP is *lower* than it would have been at the steady-state cutoff. Output falls relative to what a higher cutoff would have delivered; consumption falls with it; and the bond Euler equation then requires R_{t+1}^B to *rise* to restore household indifference between consuming today and tomorrow. The higher risk-free rate tightens the lending threshold, pushing $\bar{\omega}_{t+1}$ *up*.
- *Period $t + 2$.* The higher $\bar{\omega}_{t+1}$ raises $E[\omega \mid \omega \geq \bar{\omega}_{t+1}]$; effective TFP is now *higher*. Output recovers; consumption rises; R_{t+2}^B *falls*; and $\bar{\omega}_{t+2}$ *falls* again.
- *Subsequent periods.* The cycle repeats with alternating sign, decaying at a rate of approximately 0.69 per period.

The root cause of the oscillation is a sign mismatch between the cutoff’s forward-looking response and the lagged selection effect it triggers. The cutoff $\bar{\omega}_t$ and the TFP factor $E[\omega \mid \omega \geq \bar{\omega}_t]$ always move in the same direction—a higher cutoff selects a more productive cohort. But the cutoff’s forward-looking response to a positive shock is to *fall* on impact, and a lower cutoff then depresses effective TFP in the following period. This delayed negative feedback is what generates the negative autoregressive coefficient in the law of motion for $\bar{\omega}$ and the corresponding oscillatory dynamics in the right column of Figure 5.

8.2 Comparison of the No-default and Baseline Models

In the baseline model, aggregate output depends on the *upper* cutoff $\bar{\omega}_{t-1}$, not on $\bar{\omega}_t$. As explained in Appendix G, the forward-looking condition that pins down the lower cutoff is unchanged and given by equation (31), but $\bar{\omega}_t$ does not enter the aggregate output equation directly in the baseline model. The two cutoffs are linked through the share of defaulting firms $\frac{\mu_{\eta_1, \eta_2}(\bar{\omega}_t) - \mu_{\eta_1, \eta_2}(\bar{\omega}_{t-1})}{1 - \mu_{\eta_1, \eta_2}(\bar{\omega}_t)}$, but they respond with a degree of independence to aggregate shocks.

Two additional stabilizing forces are present in the baseline model that are absent in the no-default model. First, in the baseline model, when aggregate conditions change, the mass of defaulting firms can expand or contract almost immediately, absorbing much of the

shock before it has the opportunity to feed back through the production cutoff. Second, the aggregate equity return equation in the baseline model contains diversion-correction terms—involving the diversion scaling factor F_{t-1} , the aggregate diversion D_{t-1} , and the diversion function evaluated at the upper cutoff $\Theta(\bar{\omega}_{t-1})$ —that create a positive feedback channel from past defaulter activity to current returns. This positive channel offsets the negative selection channel described above, with the net result that the dominant eigenvalue for $\bar{\omega}_t$ remains positive. The oscillatory mode is suppressed, and the IRFs in the left column of Figure 5 decay smoothly toward zero.

8.3 How Much TFP Variance is Endogenous?

The impulse responses in Figure 5 are informative about the dynamic structure of the endogenous TFP component, but they are computed shock by shock. To obtain an unconditional, model-wide assessment we simulate each model and ask what fraction of the variance of the Solow residual is attributable to the exogenous process Z_t alone.

The calculation proceeds as follows. We draw 50 simulated samples of 500 quarters each from the second-order perturbation solution of each model, using all three shocks simultaneously. All series are HP-filtered with smoothing parameter $\lambda = 1,600$ before moments are computed. Because production in period t uses the cutoff determined in period $t - 1$ as shown in equation (G.20), the measured Solow residual decomposes as

$$\log \text{TFP}_t = \log \Phi_t + \log Z_t, \quad (45)$$

where Φ_t is the endogenous TFP prefactor from equation (42), which already uses the lagged cutoff $\bar{\omega}_{t-1}$. The variance of $\log \text{TFP}_t$ therefore satisfies

$$\text{Var}(\log \text{TFP}_t) = \text{Var}(\log Z_t) + \text{Var}(\log \Phi_t) + 2 \text{Cov}(\log \Phi_t, \log Z_t). \quad (46)$$

The headline statistic is the ratio $\text{Var}(\log Z_t)/\text{Var}(\log \text{TFP}_t)$, which equals one in a frictionless RBC economy where Φ_t is constant, and falls below one whenever the endogenous component adds variance to the Solow residual. The gap relative to the canonical RBC

model,

$$1 - \frac{\text{Var}(\log Z_t)}{\text{Var}(\log \text{TFP}_t)}, \quad (47)$$

measures the share of Solow TFP variance that cannot be accounted for by the exogenous process alone. Although this gap combines $\text{Var}(\log \Phi_t)$ with its covariance with $\log Z_t$, both terms are properly attributed to the endogenous mechanism. Were the reallocation channel shut down, Φ_t would be constant and *both* $\text{Var}(\log \Phi_t)$ and $\text{Cov}(\log \Phi_t, \log Z_t)$ would be zero; the covariance is itself a product of the endogenous response of Φ_t to the shocks. The entire gap therefore measures the contribution of the endogenous channel, including the way it co-moves with the exogenous process.

Table 7 reports results averaged across the 50 simulations. In the baseline model, roughly 30 percent of the variance of the Solow TFP cannot be explained by exogenous productivity fluctuations alone; this 30 percent is accounted for by the endogenous mechanism. The no-default model also amplifies Solow TFP variance beyond what Z_t alone can generate, but to a lesser extent: the gap is roughly 20 percent, about three-quarters of the baseline figure. The difference between the two models reflects the additional default margin present in the baseline. When conditions change, the mass of defaulting firms adjusts rapidly, concentrating production in the most productive firms more aggressively than the single-cutoff mechanism can, and thereby generating a larger endogenous component of measured TFP.

9 Conclusion

We have developed a model in which informational frictions give rise to credit misallocation. In our model, as in Stiglitz and Weiss (1981), credit is rationed.

Each firm is endowed with two technologies: a production technology and a financial intermediation technology. The efficiency of the production technology is an idiosyncratic draw that is private information. Because households cannot observe this draw, they end up financing all firms equally. Depending on the realization of their production-technology draw, firms optimally choose which technology to deploy: Those with sufficiently low draws act as financial intermediaries and lend to firms with high draws. Borrowers are subject to a default decision.

We show that this tractable model can capture some key facets of the data. First and foremost, it is able to capture the average within-industry dispersion in productivity observed at the plant level in the United States. The model also reproduces the signs of a set of untargeted partial correlations: within-industry productivity dispersion narrows when GDP growth is above average, and business-loan delinquencies rise when growth is depressed and when real interest rates are high.

In our model, changes in productivity dispersion are linked to economic conditions, implying that a component of aggregate TFP evolves endogenously. We have shown that this endogenous component of productivity can induce notable TFP movements for realistically sized shocks.

We have further isolated the specific contribution of strategic default to the endogenous TFP component by contrasting the baseline model with a no-default variant in which the two-cutoff structure collapses to a single lending threshold. Unconditional simulations show that roughly 30 percent of the variance of the Solow residual is endogenously generated in the baseline model, compared with roughly 20 percent in the no-default variant. Strategic default therefore accounts for a meaningful share of the endogenous TFP component—acting not merely as an amplifier that concentrates production among high-productivity firms, but also as a stabiliser of the endogenous component.

We hope that extensions of our tractable model can serve as a building block in the exploration of how alternative policies affect productivity. Extensions could allow the study of fiscal and monetary policy choices. Going beyond our current analysis, we can speculate that expansionary fiscal policies could have additional desirable effects in our model as the associated increases in real interest rates would force some unproductive firms to quit. By contrast, monetary policy would face additional challenges in our setup. The typical New Keynesian rationale to lower policy rates in a downturn of aggregate demand could inefficiently keep low-productivity firms afloat. In sum, we see our work as opening promising avenues for further research.

References

- Aghion, P., & Howitt, P. (1992). A Model of Growth Through Creative Destruction. *Econometrica*, 60(2), 323–351.
- Arellano, C., Bai, Y., & Kehoe, P. J. (2019). Financial Frictions and Fluctuations in Volatility. *Journal of Political Economy*, 127(5), 2049–2103.
- Barlevy, G. (2002). The Sullyng Effect of Recessions. *The Review of Economic Studies*, 69(1), 65–96.
- Bernanke, B., & Gertler, M. (1990). Financial fragility and economic performance. *The quarterly journal of economics*, 105(1), 87–114.
- Bernanke, B. S., Gertler, M., & Gilchrist, S. (1999). The financial accelerator in a quantitative business cycle framework. In J. B. Taylor & M. Woodford (Eds.), *Handbook of Macroeconomics* (pp. 1341–1393). Elsevier.
- Bernard, A. B., Eaton, J., Jensen, J. B., & Kortum, S. (2003). Plants and Productivity in International Trade. *American Economic Review*, 93(4), 1268–1290.
- Buera, F. J., & Moll, B. (2015). Aggregate Implications of a Credit Crunch: The Importance of Heterogeneity. *American Economic Journal: Macroeconomics*, 7(3), 1–42.
- Caballero, R. J., & Hammour, M. L. (1994). The Cleansing Effect of Recessions. *American Economic Review*, 84(5), 1350–1368.
- Caballero, R. J., & Hammour, M. L. (1996). On the Timing and Efficiency of Creative Destruction. *The Quarterly Journal of Economics*, 111(3), 805–852.
- Caballero, R. J., & Hammour, M. L. (1998). The Macroeconomics of Specificity. *Journal of Political Economy*, 106(4), 724–767.
- Carlstrom, C. T., & Fuerst, T. S. (1997). Agency Costs, Net Worth, and Business Fluctuations: A Computable General Equilibrium Analysis. *American Economic Review*, 87(5), 893–910.
- Christiano, L. J., Motto, R., & Rostagno, M. (2014). Risk Shocks. *American Economic Review*, 104(1), 27–65.
- Clementi, G. L., & Hopenhayn, H. A. (2006). A Theory of Financing Constraints and Firm Dynamics. *Quarterly Journal of Economics*, 121(1), 229–265.

- Comin, D., & Gertler, M. (2006). Medium-Term Business Cycles. *American Economic Review*, *96*(3), 523–551.
- Cooley, T. F., & Quadrini, V. (2001). Financial Markets and Firm Dynamics. *American Economic Review*, *91*(5), 1286–1310.
- Cunningham, C., Foster, L., Grim, C., Haltiwanger, J., Pablonia, S. W., Stewart, J., & Wolf, Z. (2023). Dispersion in dispersion: Measuring establishment-level differences in productivity. *Review of Income and Wealth*, *69*(4), 999–1032.
- Dang, J., Gillman, M., & Kejak, M. (2012). *Real Business Cycles with a Human Capital Investment Sector and Endogenous Growth: Persistence, Volatility and Labor Puzzles* (Working Paper). Cardiff Business School.
- David, J. M., & Venkateswaran, V. (2019). The Sources of Capital Misallocation. *American Economic Review*, *109*(7), 2531–2567.
- Dong, D., Liu, Z., & Wang, P. (2025). Turbulent business cycles. *Journal of Monetary Economics*, 103814.
- Foster, L., Haltiwanger, J., & Syverson, C. (2008). Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability? *American Economic Review*, *98*(1), 394–425.
- Gilchrist, S., Sim, J. W., & Zakrajšek, E. (2014, April). *Uncertainty, Financial Frictions, and Investment Dynamics* (NBER Working Papers No. 20038). National Bureau of Economic Research, Inc.
- Gomes, J. F., & Schmid, L. (2021). Equilibrium Asset Pricing with Leverage and Default. *Journal of Finance*, *76*(2), 977–1018.
- Gopinath, G., Kalemli-Özcan, Ş., Karabarbounis, L., & Villegas-Sanchez, C. (2017). Capital Allocation and Productivity in South Europe. *Quarterly Journal of Economics*, *132*(4), 1915–1967.
- Gornemann, N., Guerrón-Quintana, P. A., & Saffie, F. (2025). Real Exchange Rates and Endogenous Productivity. *American Economic Journal: Macroeconomics*, *17*(4), 204–261.

- Guo, X., Ottonello, P., Winberry, T., & Whited, T. (2025). *Firm heterogeneity and adverse selection in external finance: Micro evidence and macro implications* (tech. rep.). National Bureau of Economic Research.
- Hsieh, C.-T., & Klenow, P. J. (2009). Misallocation and Manufacturing TFP in China and India. *The Quarterly Journal of Economics*, *124*(4), 1403–1448.
- Jo, I. H., Khan, A., Senga, T., & Thomas, J. K. (2021). *Firm debt and default over the pandemic and recovery* (tech. rep.). Mimeo, Ohio State University.
- Jones, L. E., Manuelli, R. E., & Siu, H. E. (2005). Fluctuations in Convex Models of Endogenous Growth, II: Business Cycle Properties. *Review of Economic Dynamics*, *8*(4), 805–828.
- Jovanovic, B. (1982). Selection and the Evolution of Industry. *Econometrica*, *50*(3), 649–670.
- Khan, A., & Thomas, J. K. (2013). Credit shocks and aggregate fluctuations in an economy with production heterogeneity. *Journal of Political Economy*, *121*(6), 1055–1107.
- Kim, J., Kim, S., Schaumburg, E., & Sims, C. A. (2008). Calculating and using second-order accurate solutions of discrete time dynamic equilibrium models. *Journal of Economic Dynamics and Control*, *32*, 3397–3414.
- King, R. G., Plosser, C. I., & Rebelo, S. T. (1988). Production, growth and business cycles : I. the basic neoclassical model. *Journal of Monetary Economics*, *21*(2-3), 195–232.
- King, R. G., & Rebelo, S. T. (1990). Public Policy and Economic Growth: Developing Neoclassical Implications. *Journal of Political Economy*, *98*(5), S126–S150.
- Liu, Z., & Wang, P. (2014). Credit Constraints and Self-Fulfilling Business Cycles. *American Economic Journal: Macroeconomics*, *6*(1), 32–69.
- Lucas, R. E. (1988). On the mechanics of economic development. *Journal of monetary economics*, *22*(1), 3–42.
- Melitz, M. J. (2003). The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica*, *71*(6), 1695–1725.
- Midrigan, V., & Xu, D. Y. (2014). Finance and Misallocation: Evidence from Plant-Level Data. *American Economic Review*, *104*(2), 422–458.
- Moll, B. (2014). Productivity Losses from Financial Frictions: Can Self-Financing Undo Capital Misallocation? *American Economic Review*, *104*(10), 3186–3221.

- Müller, U. K., & Watson, M. W. (2008). Testing Models of Low-Frequency Variability. *Econometrica*, 76(5), 979–1016.
- Ottonello, P., & Winberry, T. (2020). Financial heterogeneity and the investment channel of monetary policy. *Econometrica*, 88(6), 2473–2502.
- Pack, H. (1994). Endogenous Growth Theory: Intellectual Appeal and Empirical Shortcomings. *Journal of Economic Perspectives*, 8(1), 55–72.
- Paulson, A. L., Townsend, R. M., & Karaivanov, A. (2006). Distinguishing limited liability from moral hazard in a model of entrepreneurship. *Journal of political Economy*, 114(1), 100–144.
- Rebelo, S. T. (1991). Long-Run Policy Analysis and Long-Run Growth. *Journal of Political Economy*, 99(3), 500–521.
- Restuccia, D., & Rogerson, R. (2008). Policy Distortions and Aggregate Productivity with Heterogeneous Establishments. *Review of Economic Dynamics*, 11(4), 707–720.
- Romer, P. M. (1990). Endogenous Technological Change. *Journal of Political Economy*, 98(5), S71–S102.
- Stiglitz, J. E., & Weiss, A. (1981). Credit rationing in markets with imperfect information. *The American Economic Review*, 71(3), 393–410.

Tables and Figures

Table 1: Summary of firms' decisions and actions within each period

	Period t	Period $t + 1$
1	Raise equity a_t	Produce, outside option matures
2	Draw productivity level $\omega \in [0, 1]$	Repay loans to other firms
3	Lend or borrow in inter-firm market	Pay households
4	Some borrowing firms take the outside option and default	
5	Purchase physical capital	

Note: The bolded entries denote actions or events specific to our model and not part of a standard RBC model.

Table 2: Model parameters

	Value	Description	
<i>Conventional</i>			
β	0.9925	Discount rate	
α	0.3	Capital share in production	
δ	0.01	Depreciation rate	
ν	2	Inverse Frisch elasticity of labor supply	
<i>Estimated to match first moments with steady-state conditions</i>			<i>Targets/Explanation</i>
η_1	1.543	First parameter of Beta distribution	} jointly to match the average 1) 90-10 within-industry productivity dispersion, 2) spread between the return on assets and the outside-option return (prime rate minus one-year Treasury), 3) share of bank credit, and 4) delinquency rate on business loans
η_2	2.735	Second parameter of Beta distribution	
ξ	0.007	Haircut on the returns from the outside option	
θ	0.0003	Average fraction of funds that can be diverted	
<i>Estimated by simulated methods of moments</i>			<i>Targets/Explanation</i>
ϕ	0.239	Investment adjustment costs	} estimated to hit the variances, correlations, and autocorrelations of 1) real output, 2) real consumption, 3) relative investment price, and 4) delinquency rate on business loans
ρ_z	0.407	Persistence of technology shock	
σ_z	0.010	s.d. of technology shock	
ρ_I	0.990	Persistence of investment technology shock	
σ_I	0.001	s.d. of investment technology shock	
ρ_ν	0.986	Persistence of consumption shock	
σ_ν	0.021	s.d. of consumption shock	
<i>Specific</i>			<i>Explanation</i>
ψ	0.5	Parameter in the function $\Theta_t(\omega) = \omega^\psi F_t$	Verifies $\Theta_t(\omega)$'s properties and the slope condition
ϑ	0.84236	Disutility of labor	Supports aggregate labor = 1 in the steady state

Table 3: SMM estimation: Data moments and model counterparts, 1987:Q1–2021:Q4

	Data	Data 5th perc.	Data 95th perc.	Model	Model 5th perc.	Model 95th perc.
Var(GDP)	1.70	0.77	2.63	2.29	1.60	3.08
Corr(GDP,Consumption)	0.91	0.25	1.57	0.84	0.76	0.90
Corr(GDP,Investment price)	0.31	0.14	0.48	0.70	0.62	0.78
Corr(GDP,Default rate)	-0.58	-0.78	-0.39	-0.59	-0.70	-0.46
Var(Consumption)	2.10	0.41	3.80	1.78	1.23	2.45
Corr(Consumption,Investment price)	0.22	0.10	0.35	0.46	0.32	0.59
Corr(Consumption,Default rate)	-0.42	-0.57	-0.27	-0.27	-0.43	-0.11
Var(Investment price)	0.56	0.40	0.72	0.62	0.47	0.78
Corr(Investment price,Default rate)	-0.48	-0.71	-0.26	-0.94	-0.96	-0.92
Var(Default rate)	0.34	0.24	0.45	0.34	0.19	0.52
Autocorr(GDP)	0.64	0.46	0.82	0.60	0.48	0.71
Autocorr(Consumption)	0.57	0.31	0.84	0.63	0.50	0.74
Autocorr(Investment price)	0.90	0.64	1.17	0.22	0.08	0.35
Autocorr(Default rate)	0.95	0.65	1.24	0.24	0.07	0.39

Note: The table reports the second moments targeted in the SMM procedure for the model calibration—variances, correlations, and autocorrelations at business-cycle frequencies. The model moments are computed from the average of 1,000 simulated samples of the same length as the observed data. Both observable and model-simulated data are HP-filtered setting the smoothing parameter to 1,600 as standard for quarterly data. The data 5th and 95th percentile columns show the bounds of a 90% confidence interval constructed under the assumption of asymptotic normality, using Newey–West HAC standard errors (Bartlett kernel with optimal bandwidth) for the sample moments. The model 5th and 95th percentile columns show the corresponding empirical percentiles from 1,000 simulated samples.

Table 4: Variance decomposition (percent)

	Output	Invest. Price	Consumption	Default Rate	Real Rate	90-10 TFP disp.
TFP	73.8	87.7	28.6	68.5	74.6	44.0
Consumption	25.1	11.5	69.4	29.0	22.9	52.9
ISP	1.1	0.9	2.0	2.5	2.6	3.2

Note: “TFP” refers to the shock to total factor productivity, Z_t . “Consumption” refers to the consumption shock, ν_{ct} . “IST” refers to the investment technology shock Z_{It} .

Table 5: Regressions of within-industry TFP dispersion (percent difference between the TFP for the 90th percentile firm and that for the 10th percentile firm) and delinquency rates, 1987–2021

	(1)	(2)	(3)	(4)
	90-10 perc.	90-10 perc.	Delinquency rate	Delinquency rate
GDP per capita, growth	-2.86 [-4.16,-1.56]	-2.72 [-3.87,-1.56]	-0.22 [-0.31,-0.12]	-0.22 [-0.31,-0.13]
Delinquency rate	-4.22 [-5.56,-2.88]	-4.38 [-5.57,-3.19]		
Lagged GDP per capita, growth		-2.31 [-3.56,-1.06]		
delinquency_rate_lag			0.84 [0.74,0.94]	0.78 [0.68,0.88]
Real interest rate				0.14 [0.04,0.24]

Note: 90% confidence intervals in brackets.

Table 6: Regressions of within-industry TFP dispersion (percent difference between the TFP for the 90th percentile firm and that for the 10th percentile firm) and delinquency rate, model-generated data

	(1)	(2)	(3)	(4)
	90-10 perc.	90-10 perc.	Delinquency rate	Delinquency rate
GDP per capita, growth	-0.70 [-1.20,-0.34]	-0.77 [-1.33,-0.33]	-0.11 [-0.19,-0.04]	0.00 [-0.03,0.03]
Delinquency rate	-4.28 [-10.47,0.04]	-4.22 [-10.59,0.09]		
Lagged GDP per capita, growth		-0.66 [-1.06,-0.25]		
Lagged delinquency rate			0.41 [0.07,0.72]	0.21 [0.02,0.44]
Real interest rate				0.26 [0.19,0.34]
<i>N</i>	34	33	34	34

90% range across replications in brackets

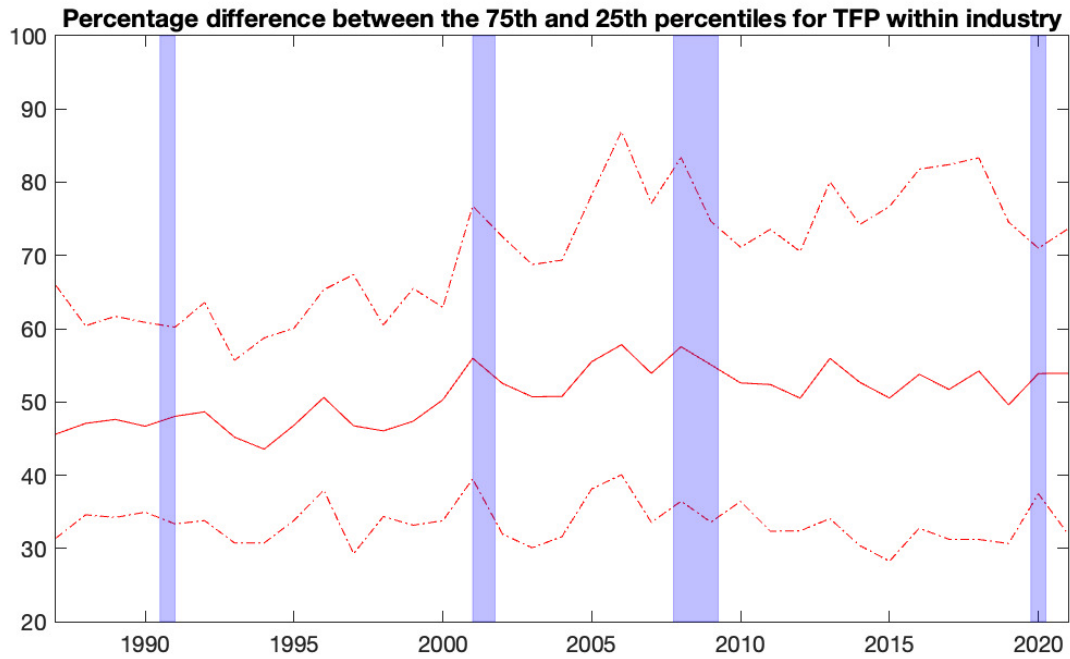

Note: The results are based on 1000 model-simulated samples of the same length as the observed data used for the counterpart regressions in Table 5. We use yearly means of quarterly model data to replicate the frequency of the observed data. 90% confidence intervals are reported in brackets. The confidence intervals for regressions on simulated data are based on the 5-95 percentiles from the 1000 data replications.

Table 7: Variance decomposition of the Solow residual

	$\text{Var}(\log Z_t)/\text{Var}(\log \text{TFP}_t)$	Gap vs. RBC (percent)
RBC benchmark	1.000	0.0
Two-cutoff model (baseline)	0.716	28.4
One-cutoff model (no strategic default)	0.788	21.2

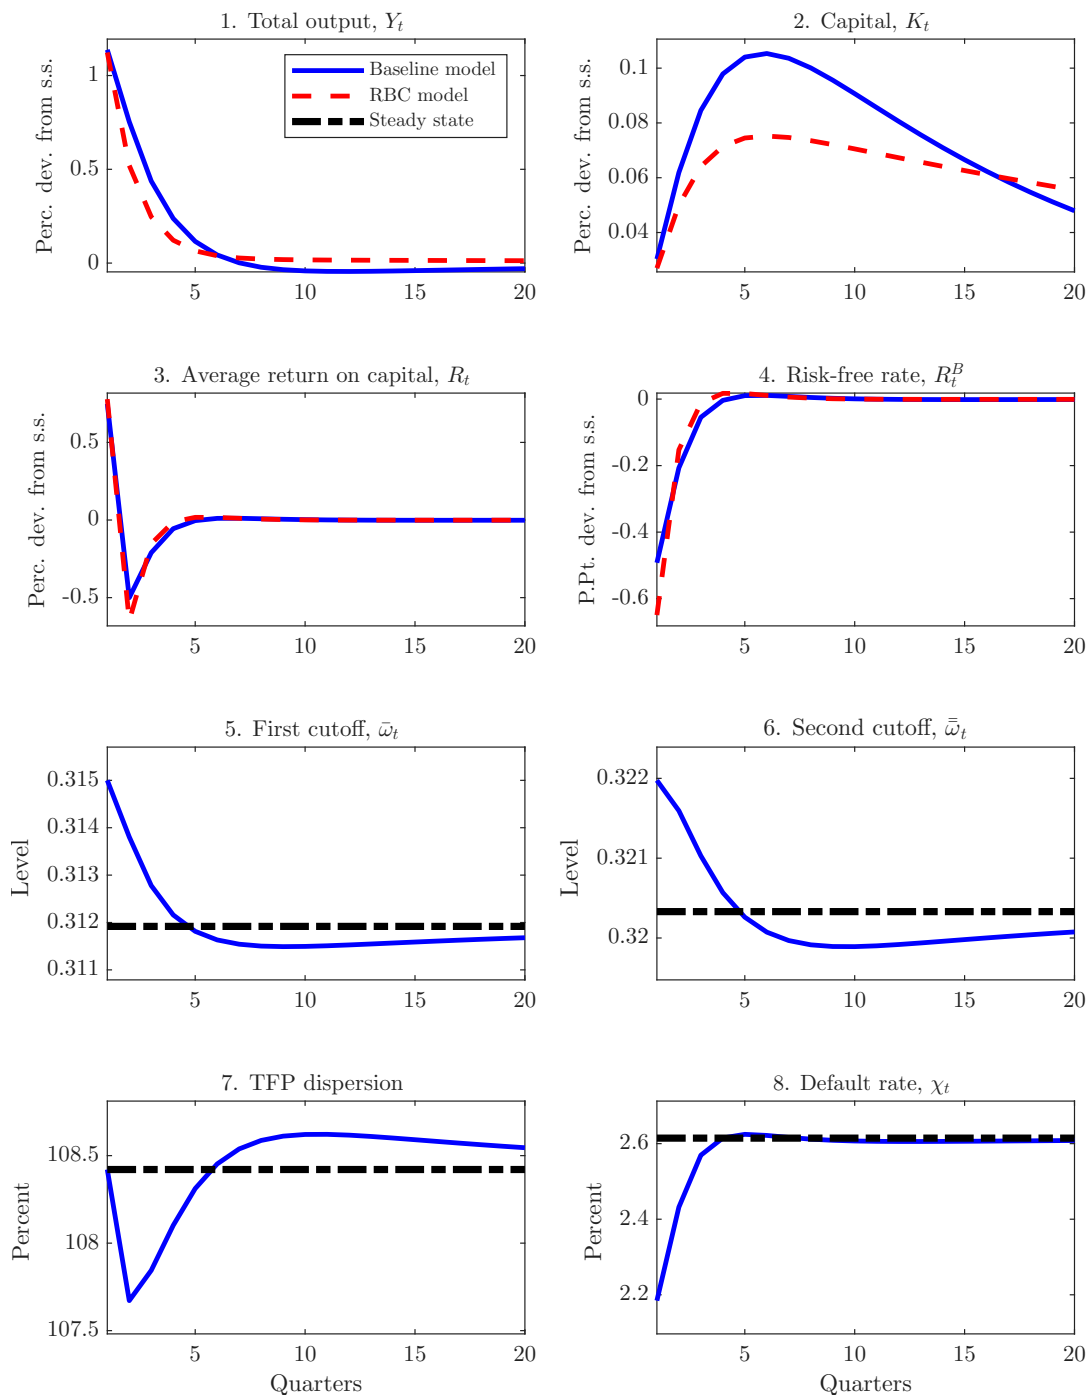
Note: HP-filtered simulations ($\lambda = 1,600$), 50 draws of 500 quarters each, all three shocks active. The ratio equals one when Φ_t is constant (RBC); the gap measures the endogenous share of Solow TFP variance.

Figure 1: TFP dispersion, 1987–2021



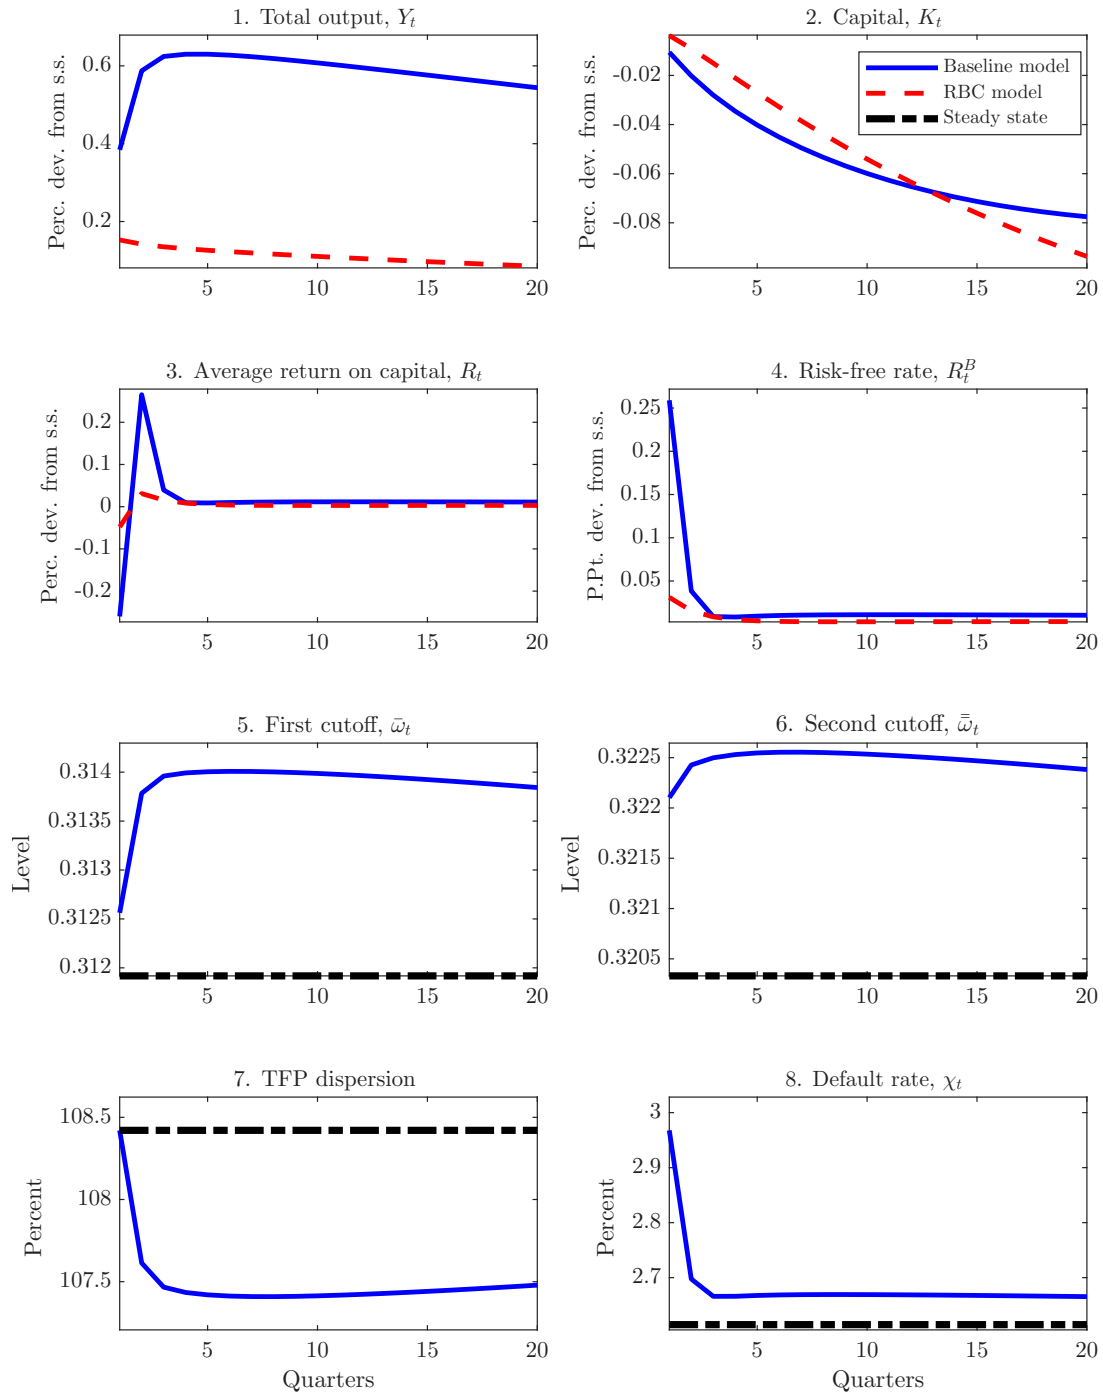
Note: Shaded areas denote recessions as dated by the NBER Business Cycle Dating Committee.
 Source: Dispersion Statistics on Productivity, a dataset published jointly by the Bureau of Labor Statistics and the Census Bureau, and authors' calculations.

Figure 2: An expansionary one-standard-deviation TFP shock



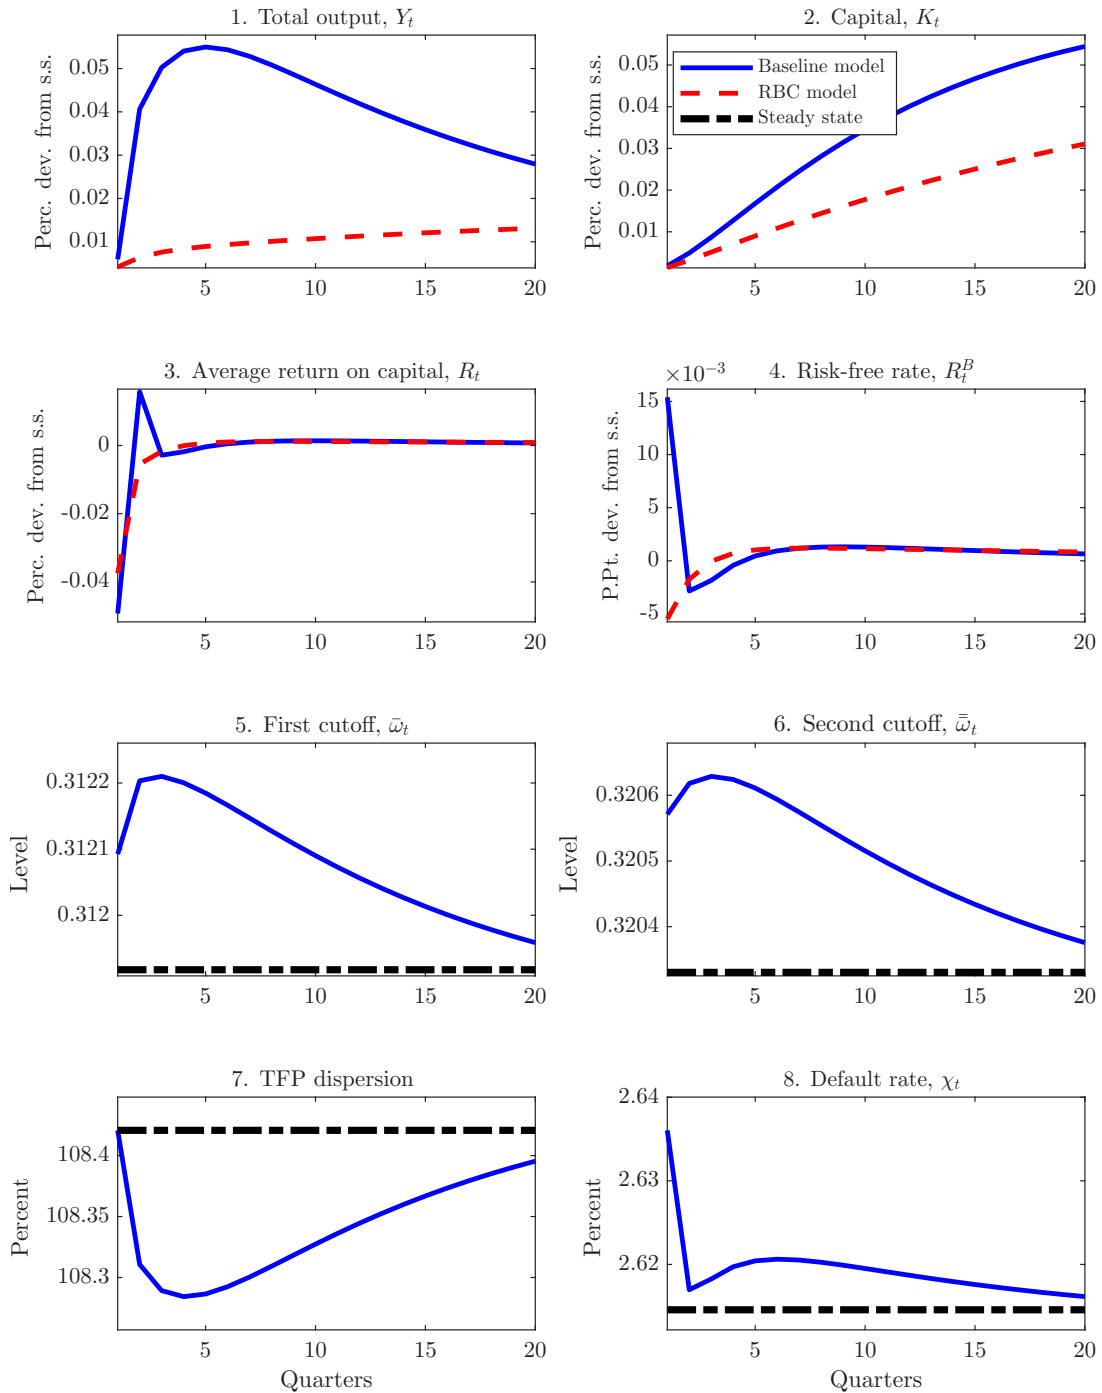
Note: “Omega bar” refers to $\bar{\omega}_t$ and “omega double bar” refers to $\bar{\bar{\omega}}_t$. The responses start from the stochastic steady state and are also shown in deviation from the stochastic steady state, where relevant.

Figure 3: An expansionary one-standard-deviation consumption shock




Note: “Omega bar” refers to $\bar{\omega}_t$ and “omega double bar” refers to $\bar{\bar{\omega}}_t$. The responses start from the stochastic steady state and are also shown in deviation from the stochastic steady state, where relevant.

Figure 4: An expansionary one-standard-deviation investment technology shock



Note: “Omega bar” refers to $\bar{\omega}_t$ and “omega double bar” refers to $\bar{\bar{\omega}}_t$. The responses start from the stochastic steady state and are also shown in deviation from the stochastic steady state, where relevant.

Figure 5: TFP decomposition: baseline vs. no-default model



Note: Each panel shows the impulse response of the Solow residual decomposed into its two components (Section 5.2). Because $\log \text{TFP}_t = \log \Phi_t + \log Z_t$, the decomposition is *additive in log-percent units*: the blue and orange areas stack to exactly total measured log TFP at each horizon. Blue area: exogenous component $\log Z_t$ (log % deviation from steady state). Orange area: endogenous component $\log \Phi_t$ (log % deviation from steady state). Rows correspond to the three shocks; columns correspond to the two models. The y -axis scale is identical within each row to facilitate direct comparison. Left column: baseline model with strategic default (two-cutoff model, $\bar{\omega}_t$ and $\bar{\bar{\omega}}_t$). Right column: no-default model (one-cutoff model, $\bar{\omega}_t$ only; see Section G of the [online appendix](#)).