

Cyclical Fluctuations, Financial Frictions, and Productivity Differences across Firms

Luca Guerrieri¹ Jinill Kim² Arsenii Mishin³

Macro seminar, George Washington University,
2026

¹ Federal Reserve Board

² Korea University & Bank of Korea


³ HSE University

The views expressed here 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 anyone else associated with the Federal Reserve System, of the Bank of Korea, or of anyone else associated with the Bank of Korea.

01

Motivation and preview

TFP dispersion across manufacturing is large and widens in downturns



U.S. manufacturing, 86 four-digit NAICS industries (Census/BLS DiSP), 1987–2021.

- 90–10 gap $\approx 100\%$: the 90th-percentile plant is $\approx 2.1\times$ as productive as the 10th.
- Pervasive — not driven by outliers.
- Countercyclical: narrows when GDP growth is high, widens when it is low.

These are the moments we ask the model to reproduce.

When dispersion moves, measured TFP moves with it

- Measured TFP — the Solow residual — is not a primitive. It combines the productivity of individual firms *and* how efficiently inputs are allocated across them.
- So if the spread of productivity across firms widens and narrows over the cycle, so does the efficiency with which the economy uses its inputs, and so does aggregate TFP.
- Capturing that link usually means tracking the whole distribution of firms — hard to embed in a business-cycle model.

Our question

How much of measured TFP over the business cycle is endogenous to credit conditions — and can we capture it in a model simple enough to estimate?

Heterogeneous firms, no distribution to track

A private draw

Each firm draws its own productivity ω , seen by nobody else, and can lend to other firms instead of producing.



Two frictions sort

Adverse selection and moral hazard split firms into **lenders, strategic defaulters**, and **producers**.



Two cutoffs, one price

The whole cross-section collapses to $\bar{\omega}$, $\bar{\omega}$ and the inter-firm loan rate.

So the mechanism adds a handful of equations to a standard RBC model, and is solved and calibrated with off-the-shelf methods.

Preview of results: credit conditions move TFP

Headline

Credit conditions make part of *measured* TFP endogenous – about **30%** of the variance of TFP over the business cycle – and **strategic default** accounts for about one third of that endogenous component.

- A technology shock is *amplified*: dispersion narrows and production concentrates in the most productive firms.
- Shocks that leave technology untouched still move measured TFP.
- In untargeted regressions the model reproduces the signs in the data: dispersion narrows when growth is high, delinquencies rise when growth is weak and real rates are high.
- Strategic default both amplifies the endogenous component and smooths its path.

Where we differ: endogenous TFP, no distribution

- **Financial frictions and amplification.** Agency costs of lending amplify aggregate shocks — but there productivity dispersion is fixed, and default means inability to repay rather than a choice.
- **Endogenous TFP through credit.** Reallocating capital across firms moves aggregate TFP — typically at the cost of carrying a firm distribution as a state variable.
- **Misallocation.** Distortions and financial frictions keep capital from its most productive uses, at a large cost in aggregate productivity.
- **Endogenous growth and innovation.** Aggregate productivity is an endogenous object there too — but permanent and technological, where ours is transitory, financial, and plays out over the business cycle.

Frictions & rationing

Endogenous TFP

Default

Misallocation

Growth

Plan for rest of presentation

1. Data — the cyclical of productivity gaps within industries.

2. Model — firms with private productivity sort into lenders, defaulters, and producers.

3. Equilibrium characterization — two cutoffs and a loan rate; measured TFP splits into an exogenous and an endogenous part.

4. Calibration and dynamics — what the model is asked to match, and how it responds to shocks.

5. Strategic default — what that margin adds.

02

The data

How we measure dispersion, and what we regress

- **Source.** Dispersion Statistics on Productivity (DiSP), from Census microdata: 86 four-digit NAICS manufacturing industries, annual, 1987–2021.
- **Measure.** Within industry j and year t , the 90–10 gap in establishment-level TFP, in percentage points.
- **Aggregation.** Value-added-weighted average across the 86 industries, $\text{gap}_t = \sum_j w_{jt} \text{gap}_{jt}$ (BEA weights).

To ask whether that gap moves with the cycle, we run

$$\text{gap}_t = \alpha + \beta_1 g_t + \beta_2 g_{t-1} + \gamma \text{delinq}_t + \varepsilon_t,$$

g_t : real GDP-per-capita growth (BEA). delinq_t : the delinquency rate on **business loans at commercial banks** (Federal Reserve Board, Charge-Off and Delinquency Rates release), averaged to annual. Both in annual percent.

→ We read the model's **lending firms as stylized banks**, so this series is the data counterpart of the model's default rate.

Dispersion is countercyclical

OLS on annual U.S. data, 1987–2021.

	90–10 gap		Delinquency rate	
	(1)	(2)	(3)	(4)
GDP per capita, growth	−2.86	−2.72	−0.22	−0.22
Lagged GDP growth		−2.31		
Delinquency rate	−4.22	−4.38		
Lagged delinquency			0.84	0.78
Real interest rate				0.14



- A one-point higher growth rate goes with a 2.86 pp narrower 90–10 gap.
- Partial correlations, not causal estimates: *untargeted* moments.

Every coefficient shown is significant at the 10% level; intervals are in the backup. Later we run the same regressions on model data.

03 The model

Where we are going: two cutoffs, three groups

The whole cross-section reduces to this picture: two *equilibrium* cutoffs, pinned down by the model, sort firms into lenders, defaulters, and producers.



Drawing ω from a Beta distribution lets us match the wide productivity dispersion documented by DiSP: its two shape parameters are set to reproduce the observed 90–10 gap.

Environment

Production technology


Efficiency $\omega \in [0, 1]$, drawn i.i.d. each period, private to the firm:

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

Intermediation technology

Identical across firms: lend to other firms at the inter-firm rate ρ_t .

Each firm owns both and, after its draw, deploys one: **produce** or **intermediate**.



Households cannot see ω , so every firm gets the same equity a_t .

Tractable additions to the RBC model

The two-period overlapping structure is a classic alternative decentralization of the RBC model, with equity contracts channelling household savings to firms. Starting from that benchmark, our model adds the **bold** steps below.

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

Adverse selection and strategic default ration credit

Two informational frictions follow from the private draw, and together they misallocate capital.

Adverse selection

Outside investors cannot observe ω , so the composition of borrowers cannot be screened.

Moral hazard


Firms with low ω find it profitable to divert borrowed funds to an outside option rather than produce – *strategic default*.

The twist relative to Stiglitz–Weiss (1981)

There, default is *involuntary*: project returns fall short. Here it is a *strategic equilibrium choice*. Stiglitz–Weiss is static and partial-equilibrium; ours is general equilibrium with capital, a goods market, and business cycles.

Credit rationing and the screening technology

With observable ω , the first best would have the single most productive firm borrow everything and produce alone. Instead:



- Lenders are not blind: a *screening technology* tells them whether a borrower could beat the outside option by producing. It keeps the least productive out of the borrowing pool, but cannot rank the firms above $\bar{\omega}_t$.
- **Two frictions, two jobs.** Adverse selection makes the contract *common*: a higher promised rate signals nothing, since a firm that plans to divert can promise anything and repay nothing. Moral hazard sets how *large* that common loan can be: lenders price in the chance the funds are diverted.


From here: one problem per segment

Given the two frictions, each firm compares what it can earn from its two technologies. We take the three cases in turn:

- a. **Firms that borrow and produce.**
- b. **Firms that borrow and default.** The outside option and the diversion technology.
- c. **Firms that lend.** The pooled return on loans to producers and to defaulters.

Then the rest of the economy — households, capital producers, government — and a first look at what the machinery buys.

Firms that borrow and produce ($\omega \geq \bar{\omega}_t$)




Capital is **equity** plus an **inter-firm loan**; at t the firm picks a plan for every draw:

$$\begin{aligned} \max_{\{k_t(\omega), h_{t+1}(\omega), b_t(\omega)\}} \quad & E_t \left[\beta \frac{\lambda_{ct+1}}{\lambda_{ct}} \left(R_{t+1}(\omega) (a_t + b_t(\omega)) - R_{t+1}^A(\omega) a_t - \rho_t b_t(\omega) \right) \right] \\ \text{s.t.} \quad & Q_t k_t(\omega) = a_t + b_t(\omega) \quad \text{for every } \omega. \end{aligned}$$

- With constant returns the firm would borrow all it can. In equilibrium credit is rationed: every borrower gets the same contract, $b_t(\omega) = b_t$, sized by *supply*.
- Nothing in the plan then depends on ω : it collapses to one triple $(\tilde{k}_t, \tilde{h}_{t+1}, \tilde{b}_t)$, common to every producer.

Across producing firms, only ω differs



Given that all producing firms are ex-ante identical, they hold the same $(\tilde{k}_t, \tilde{h}_{t+1})$, so output *and* the return differ only through the ω draw:

$$y_{t+1}(\omega) = Z_{t+1} \omega \tilde{k}_t^\alpha \tilde{h}_{t+1}^{1-\alpha},$$
$$R_{t+1}(\omega) = \frac{\kappa_{t+1}}{Q_t} \omega + \frac{(1-\delta)Q_{t+1}}{Q_t}, \quad \kappa_{t+1} \equiv \alpha Z_{t+1} (H_{t+1}/K_t)^{1-\alpha}.$$

κ_{t+1} is *common to all producers* and positive, so the return is strictly increasing in ω .

- So ω is the only firm-specific term – what sorts firms into segments, and what lets the cross-section aggregate.

Firms that borrow and default ($\bar{\omega}_t < \omega < \bar{\bar{\omega}}_t$)



A borrower can walk away, putting what it keeps into the government bond:

$$\max_{b_t(\omega)|\omega} E_t \left[\beta^{\frac{\lambda_{ct+1}}{\lambda_{ct}}} \left\{ \underbrace{(R_t^B - \xi)}_{\text{bond, after the haircut}} \left(a_t + \underbrace{\Theta_t(\omega)}_{\text{divertible share}} b_t \right) - R_{t+1}^A(\omega) a_t \right\} \right]$$

The diversion technology

$\Theta_t(\omega)$ is non-negative, increasing and concave; the baseline takes $\Theta_t(\omega) = \omega^\psi F_t$, with F_t set so the average diverted share equals θ .

Why it is strategic

Default is a choice, not an inability to repay – and creditors cannot reach the funds already diverted.

Firms that lend ($\omega \leq \bar{\omega}_t$)



$$\max_{l_t(\omega) \leq a_t(\omega)} E_t \left[\beta \frac{\lambda_{ct+1}}{\lambda_{ct}} \left(R_t^{\text{pool}} l_t(\omega) - R_{t+1}^A(\omega) a_t(\omega) \right) \right],$$

$$R_t^{\text{pool}} = \underbrace{\rho_t \frac{1 - \mu(\bar{\bar{\omega}}_t)}{1 - \mu(\bar{\omega}_t)}}_{\text{repaid in full by producers}} + \underbrace{(1 - \theta) \frac{\mu(\bar{\bar{\omega}}_t) - \mu(\bar{\omega}_t)}{1 - \mu(\bar{\omega}_t)}}_{\text{recovered from defaulters}}$$

Why the return is pooled

Contracts cannot condition on ω , so a lender cannot tell producers from defaulters in advance.

Lending rate, ρ_t , endogenous

Set by indifference between lending and borrowing-to-default; revenue rises with the amount lent, so the constraint binds.

The rest of the economy is standard

Everything outside the firm block is textbook. Three pieces come back later:

Households

One representative household with King–Plosser–Rebelo preferences, holding equity in goods firms.

→ supplies the discount factor every firm uses, and its bond Euler equation pins down R_t^B .

Capital producers

Competitive firms build capital with quadratic investment adjustment costs (ϕ), giving a Tobin's q price Q_t .


→ the expected resale value Q_{t+1} enters every producing firm's return.

Government

Issues bonds *only to firms*, and sets the lump-sum transfer T_t so the budget constraint holds period by period.

→ the bond is the diverters' outside option; the haircut is rebated, $\Xi_t = \xi D_{t-1}$, so the friction destroys no resources.

Both kinds of shocks move measured TFP



- **Left.** Output rises more than in the frictionless model, and stays higher for longer.
- **Middle.** That extra response is the endogenous Φ_t stacking on top of the shock.
- **Right.** After a preference shock Z_t is unchanged, yet TFP still moves: all of it endogenous.

Next: what pins the two cutoffs down, and how they aggregate into Φ_t .

04

Characterizing the equilibrium

Posit a solution, then verify no firm wants to switch

Our strategy is to posit a solution that satisfies the firms' first-order conditions segment by segment, and then verify that no firm *can* or *has an incentive* to switch to a different segment.

1. Posit

Firms sort by ω : lend below $\bar{\omega}_t$, borrow and default in between, borrow and produce above $\bar{\bar{\omega}}_t$.

2. Pin down

Indifference at each cutoff, plus the lenders' zero-profit condition, give three conditions for ρ_t , $\bar{\omega}_t$ and $\bar{\bar{\omega}}_t$.

3. Verify

Three propositions confirm the sort: the least productive lend, the cutoffs are ordered, and no firm has an incentive to switch segment.

Proofs are in the paper's appendix; here we give the statement and the intuition.

The lower cutoff and the loan rate

At $\bar{\omega}_t$ a firm is indifferent between producing and the outside option:

$\bar{\omega}_t$ ρ_t $\bar{\bar{\omega}}_t$ Φ_t

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

This is the test the *screening technology* lets lenders apply: can the borrower beat the outside option by producing? Since the return rises with ω , every firm below $\bar{\omega}_t$ fails it.

The loan rate ρ_t makes that same firm indifferent between lending and borrowing to default:

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

What a firm earns by lending does not depend on its own ω , while what it earns by diverting rises with ω . So below $\bar{\omega}_t$ lending wins, and above it borrowing does.

The upper cutoff

At $\bar{\bar{\omega}}_t$ the marginal firm is indifferent between diverting and producing:

$\bar{\omega}_t$ ρ_t $\bar{\bar{\omega}}_t$ Φ_t

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

Above $\bar{\bar{\omega}}_t$ producing wins, below it diverting does – which is what makes $\bar{\bar{\omega}}_t$ the production margin, the one that enters aggregate output.

Three propositions verify the posited sort

Each rests on a property of the diversion technology $\Theta_t(\omega)$:

1. The least productive lend

With Θ_t non-negative and *increasing*, no firm below $\bar{\omega}_t$ wants to deviate from lending.

2. The cutoffs are ordered

With Θ_t non-negative, $\bar{\bar{\omega}}_t \geq \bar{\omega}_t$: the three segments are always well defined.

3. The sort holds

Adding *concavity* and the slope condition, firms sort as posited: lend, divert, produce.

Proposition 3 needs one more ingredient – a *slope condition*, evaluated at $\bar{\omega}_t$:

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

Intuitively, it guarantees that the returns from defaulting and from producing cross exactly once, at $\bar{\bar{\omega}}_t$. Proofs are in the paper's appendix.

Measured TFP = endogenous Φ $\times$ exogenous Z

$\bar{\omega}_t$ ρ_t $\bar{\bar{\omega}}_t$ Φ_t

$$Y_t = \Phi_t Z_t K_{t-1}^\alpha H_t^{1-\alpha}, \quad \text{TFP}_t = \Phi_t \cdot Z_t,$$
$$\Phi_t = c_\alpha \frac{1 - \mu_{\eta_1+1, \eta_2}(\bar{\bar{\omega}}_{t-1})}{1 - \mu_{\eta_1, \eta_2}(\bar{\bar{\omega}}_{t-1})} = E[\omega \mid \omega \geq \bar{\bar{\omega}}_{t-1}].$$

- Measured TFP is Z_t times the *average productivity of producing firms*.
- $\bar{\bar{\omega}}$ responds to financial conditions, so Φ_t moves with Z_t fixed.
- Note the timing: production at t uses the cutoff set at $t-1$.
- RBC limit: $\bar{\bar{\omega}} \rightarrow 1 \Rightarrow \Phi_t \rightarrow 1$, and TFP tracks Z_t one-for-one.

$\mu_{a,b}$: Beta(a, b) CDF; $c_\alpha = E[\omega] = \eta_1/(\eta_1 + \eta_2)$.

05

Calibration and solution

Parameters come from three places

Conventional

Standard values from the literature – for instance $\beta = 0.9925$ and $\alpha = 0.30$.

KPR preferences impose log utility in consumption.

Steady-state targets

η_1 , η_2 , the haircut ξ , the divertible share θ .

Matched jointly to four first moments.

Estimated by SMM

The three shock processes and the adjustment cost ϕ .

Second moments, 1987:Q1–2021:Q4.

Four steady-state targets, four parameters

Targets (1987–2021 averages)

- 90–10 dispersion: $2.08\times$
- prime–Treasury spread
- bank-credit share, $\approx 47\%$
- delinquency rate, 2.6%


jointly

Parameters

- $\eta_1 = 1.543, \eta_2 = 2.735$
(shape of the ω distribution)
- $\xi = 0.007$ (haircut)
- $\theta = 0.0003$ (divertible share)

All four targets depend on the shape of the distribution, so the four parameters are matched jointly rather than one to one. The model's 90–10 gap is a truncated-Beta quantile ratio over producing firms.

Fourteen cyclical moments are targeted

Four series

- real GDP
- real consumption
- relative price of investment
- business-loan delinquency rate



Fourteen moments

- 4 variances
- 6 cross-correlations
- 4 first-order autocorrelations

Data and model are filtered the same way, to focus on variation over the business cycle.

The model matches the cyclical moments targeted

14 of 14

signs matched

10 of 14

data and model 90% intervals overlap

- The four misses: the correlation of the investment price with GDP and with the delinquency rate, and the autocorrelations of those same two series.
- Three involve the investment price: an asset price in the model, tied to investment growth by construction; a goods price in the data, driven largely by equipment technology.

Which shocks drive what

Second order, not first

What we use

An exact second-order perturbation solution, with the pruning algorithm of Kim, Kim, Schaumburg and Sims.

Why not first order

In simulated samples it let the upper cutoff $\bar{\bar{\omega}}_t$ fall below the lower cutoff $\bar{\omega}_t$ – a theoretical impossibility.

- We check the second-order solution against a fully nonlinear perfect-foresight path: it tracks it closely, while the first-order solution visibly departs.

Accuracy check


06

Dynamics and validation

Plan for this section


1. **A technology shock, step by step** – what happens to the return on capital, then to the cutoffs, then to default, dispersion and output.
2. **Shocks that leave technology untouched** – preference and investment-specific shocks move TFP through the real rate alone.
3. **Untargeted moments** – the same regressions, run on model data.

The technology shock fades faster than the capital it builds




A one-standard-deviation shock raises Z by 1%, but with $\rho_Z = 0.41$ only 0.4% of that is left a quarter later. The investment it sets off bids the price of capital up on impact and lets it fall back: the capital gain is +0.7 in the first quarter and -0.5 in the second. The average return traces the same path – it is the resale term, not the marginal product, that turns the return negative from period 2.

Both cutoffs rise, and the band between them narrows



The expected return to producing falls, so both cutoffs *rise*: the marginal borrower tips toward lending, the marginal producer toward diverting. They do not rise equally: the marginal producer is *levered*, running $a + b$ of capital on a of equity, so it absorbs the same change in returns with a smaller move in $\bar{\omega}$ than the unlevered lender needs in $\bar{\omega}$. The band narrows. A higher $\bar{\omega}$ also compresses the producing pool, so dispersion narrows — but only from period 2, since production at t uses the cutoff set at $t - 1$.

Fewer defaulters, a tighter loan spread




The band between the cutoffs is the mass of firms that divert, and lenders price exactly that mass into the loan rate: the spread over the risk-free rate is compensation for expected default. The two left panels therefore trace the same path. Both moves are small because the divertible share is only $\theta = 0.03\%$, so a default costs lenders little. Output rises on impact through the standard channel and stays above steady state as the leaner producing pool lifts measured TFP.

Walking through the technology shock

- a. On impact, higher Z_t raises the marginal product of capital and lifts output – the standard channel.
- b. The shock decays, so from period 2 the *expected* return on capital falls, mostly through the expected resale price of capital.
- c. A lower expected return pushes the production cutoff $\bar{\omega}_t$ *up*: only more productive firms still clear it.
- d. Dispersion among producers narrows and Φ_t rises. The endogenous component switches on with a one-period lag.
- e. The outside-option return R_t^B falls as well, but by less, so $\bar{\omega}_t$ rises too – by more than $\bar{\bar{\omega}}_t$. The default rate *falls*: defaults are countercyclical.

The two cutoffs play different roles: $\bar{\bar{\omega}}_{t-1}$ enters production directly, while $\bar{\omega}_t$ is purely forward-looking.

Endogenous Φ adds amplification and persistence




Measured TFP = $\log Z_t + \log \Phi_t$.

- Blue: exogenous Z_t .
- Orange: endogenous Φ_t , switching on from period 2.
- The endogenous component stacks *on top* of the shock – amplification and added persistence.

All three shocks



A preference shock squeezes producers



An expansionary preference shock raises the real rate. That makes the outside option — divert the loan into the bond — pay more, so the bar for producing rises. The firms just above the old bar drop out; the ones left are better able to withstand the squeeze. The producing pool is leaner and dispersion falls.

All 8 panels

Non-technology shocks move endogenous TFP too



Both shocks leave Z_t *exactly* unchanged, so every bit of the TFP response is endogenous. Note the squeeze on producers can come from either side: here the outside option improves, while after a technology shock the expected return to producing falls. Either way the marginal producers drop out and Φ_t rises.

The same regressions, now on model data

The two regressions we ran on the data, with the model's counterparts of every series:

$$\begin{aligned}\text{gap}_t &= \alpha + \beta_1 g_t + \beta_2 g_{t-1} + \gamma \text{delinq}_t + \varepsilon_t, \\ \text{delinq}_t &= a + \rho \text{delinq}_{t-1} + b g_t + c r_t + u_t.\end{aligned}$$

The impulse responses already tell us what signs to expect:

- $\beta_1 < 0$: the technology shock raises output while it raises $\bar{\omega}$, so the 90–10 gap narrows as growth picks up.
- $b < 0$: the same shock thins the mass of defaulters — default is countercyclical.
- $c > 0$: preference and investment-specific shocks push the real rate and the default rate up together.

Same OLS as on the data, run on 1,000 simulated samples aggregated to annual frequency. None of these moments is targeted.

Untargeted moments: the model matches the signs

	90–10 dispersion		Delinquency rate	
	Data	Model	Data	Model
GDP-per-capita growth	−2.86	−0.70	−0.22	−0.11
Delinquency rate	−4.22	−4.28		
Lagged delinquency			0.84	0.41
Real interest rate (spec. 4)			0.14	0.26

- Every model coefficient matches the sign in the data. Shading of the model entry: **green** = magnitude essentially identical to the data, **amber** = same order of magnitude.
- Same OLS run on data and on 1,000 simulated samples; *not* targeted in calibration.

Reading the regressions through the model


- **Dispersion and growth.** The same technology shock that raises output narrows the 90–10 gap, which is the negative coefficient in columns 1–2.
- **Dispersion and delinquencies.** Both are driven by the cutoffs, so they co-move; the model's point estimate (-4.2) is essentially the one in the data, though its range is wide.
- **Delinquencies, growth and rates.** Preference and investment shocks push the real rate and the default rate up together, while the technology shock raises growth and pulls delinquencies down.
- **One cell does not line up.** With the real rate included, the model's growth coefficient is essentially zero: the real rate absorbs the cyclical variation that growth carries on its own.

07

What strategic default adds

To see what strategic default brings, take it away

A distinctive feature of the model is **strategic default**: borrowers choose to divert rather than repay, and the mass that does so moves with aggregate conditions.



Parameters are held at their baseline values across the two models, so every difference comes from the default margin itself.

Strategic default amplifies and stabilizes

Strategic default serves a dual purpose: it **amplifies** the endogenous TFP component by concentrating production among high-productivity firms, and it **stabilizes** that component's dynamic path.

Amplifies

When conditions change, the mass of defaulting firms adjusts rapidly — more aggressively than a single cutoff can. Worth about a third of the endogenous component.

Stabilizes



Firms near the default threshold adjust first, spreading the reallocation across two margins. With one margin, the endogenous component oscillates.

And it is observable

Default is an equilibrium choice, not a knife-edge case ruled out by contract design — and its counterpart, the delinquency rate, is targeted in the estimation.

The rest of this segment takes these in turn, against the no-default variant.


Two margins are smoother than one



Solid: baseline (smooth, persistent). Dashed: no default (oscillatory). With strategic default the mass of defaulters adjusts almost immediately, absorbing much of the shock before it feeds back through the production cutoff; diversion-correction terms in the aggregate equity return add positive feedback, so the dominant root for $\bar{\omega}_t$ stays positive and the oscillatory mode is suppressed.

About 30% of TFP variance is endogenous

≈ 30% of business-cycle TFP variance is endogenous



Gap vs. RBC = $1 - \frac{\text{Var}(\log Z_t)}{\text{Var}(\log \text{TFP}_t)}$, all three shocks active.

- $\Rightarrow$ strategic default accounts for about one third of the endogenous component.

Decomposition table

08

Conclusion

Takeaways

Credit conditions make about **30%** of the variance of TFP over the business cycle **endogenous**, and about **one third** of that comes from **strategic default**.

A tractable model

Heterogeneous firms, yet two cutoffs and a loan rate — no distribution to track.

Strategic default does two jobs

It amplifies the endogenous component, and it smooths the dynamics.

Fit to untargeted moments

Correlations of dispersion, growth, and delinquencies that the calibration never targeted.

A building block for policy analysis

We see the model as a reusable building block for studying how policy affects productivity. Two directions, both speculative:

Fiscal policy

Expansionary fiscal policy could have additional desirable effects here: the associated rise in real interest rates would force some unproductive firms to quit.

Monetary policy

Monetary policy would face an additional challenge: the usual New Keynesian case for cutting rates in a downturn could inefficiently keep low-productivity firms afloat.

Both go beyond the current analysis — the model here has no fiscal or monetary block.

Thank you

Luca Guerrieri, Jinill Kim, Arsenii Mishin

Endogenous TFP & Financial Frictions

Backup slides follow.

Backup – how DiSP measures productivity

Data. Census Bureau microdata – the Annual Survey of Manufactures and the Census of Manufactures, linked through the Longitudinal Business Database. DiSP is an experimental product built jointly by BLS and the Census Bureau (Cunningham et al., 2023).

Output. Real output is the value of shipments, adjusted for inventory changes and resales, deflated by six-digit NAICS shipment deflators.

TFP. Gross-output based, cost-share weighted, Cobb–Douglas in logs:

$$\ln \text{TFP}_{et} = \ln Q_{et} - \alpha_K \ln K_{et} - \alpha_L \ln \text{TH}_{et} - \alpha_M \ln M_{et} - \alpha_E \ln E_{et},$$

with hours worked TH, productive capital K , materials M and energy E ; the elasticities are each six-digit industry's expenditure shares in total cost.

Their own finding. Across 1997–2016, within-industry TFP dispersion is significantly countercyclical.

[◀ Back](#)

Backup – the delinquency series

Source. Delinquency rate on business loans at commercial banks, Federal Reserve Board Charge-Off and Delinquency Rates release; quarterly, averaged within the year for the annual regressions.

Definition. The Board counts a loan as delinquent when it is past due thirty days or more and still accruing interest, or in nonaccrual status; the rate is delinquent loans over total business loans outstanding.

Model counterpart. χ_t , the share of borrowing firms that divert. The model does not separate delinquency, charge-offs and outright default: a diverting firm fails to repay in full, which we treat as a single default event.

Lending firms as stylized banks. All external funds are intermediated through lending firms, so L_t/B_t^{tot} is the model analogue of the share of firm financing provided by the banking system – about 47% of US private credit.

Where it is used. Sample mean of 2.6% is a steady-state target for χ ; the bank-credit share is a second target; the prime-rate–Treasury spread pins ξ . The delinquency *moments* enter the SMM objective, and the delinquency *regressions* are untargeted.

Backup – how the dispersion series is built

Source. Dispersion Statistics on Productivity (DiSP), an experimental release built from Census Bureau microdata: 86 four-digit NAICS manufacturing industries, annual, 1987–2021. We use the second-moment measure of establishment-level total factor productivity.

Within-industry gap. For each industry and year, the difference in establishment TFP between the 90th and the 10th percentile, in percentage points.

Weights. The cross-industry average uses each industry's share of value added, from the Bureau of Economic Analysis. Several NAICS manufacturing industries map to the same BLS code, so their value added is aggregated before the weights are formed.

Cross-industry spread. The bands in the figure are the industries at the 15th and 85th percentiles of the gap in each year – the average is not driven by a few outliers.

Calibration target. The sample mean of the weighted average gap puts the 90th-percentile plant at $2.08\times$ the 10th.

Backup – related literature (1/6)

Credit frictions & amplification (costly state verification).

- Carlstrom–Fuerst (1997); Bernanke–Gertler–Gilchrist (1999); Christiano–Motto–Rostagno (2014): agency costs of intermediation amplify and propagate shocks.
- But productivity dispersion is *static / exogenous*, and default is *not strategic* (inability to repay). Heterogeneity is project-return risk, not a producing-firm selection margin.

Credit rationing.

- Stiglitz–Weiss (1981): the interest rate screens applicants (adverse selection) and shifts risk-taking (incentives). We add a *strategic-diversion* margin inside general equilibrium.

Backup – related literature (2/6)

Closest on the endogenous-TFP channel.

- Khan–Thomas (2013): collateral constraints + partial irreversibility generate persistent endogenous TFP, with equilibrium default ruled out. We obtain the same channel more parsimoniously: two cutoffs and a price, not a high-dimensional firm distribution.
- Buera–Moll (2015): under log utility, a tightening of collateral constraints is isomorphic to an exogenous TFP shock, so the endogenous component is not separately identified. Ours is, even in the i.i.d. limit.
- Liu–Wang (2014); Dong–Liu–Wang (2025): borrowing limits are exogenous and no firm defaults in equilibrium. Here the endogenous mass of defaulting firms is central to the TFP response.
- Moll (2014): with persistent productivity, self-financing erodes misallocation; with i.i.d. shocks TFP is driven by exogenous factors. We show endogenous TFP dynamics arise *even with i.i.d.* productivity.

Backup – related literature (3/6)

Heterogeneous firms, financial frictions & default.

- Gilchrist–Sim–Zakrajsek (2014); Arellano–Bai–Kehoe (2019); Gomes–Schmid (2021).
- Jo–Khan–Senga–Thomas (2021); Ottonello–Winberry (2020); Guo–Ottonello–Winberry–Whited (2025): focus on size, leverage, and equity-financing distributions.

Closest on the default friction.

- Cao–Lorenzoni–Walentin (2019): limited enforcement – a defaulting entrepreneur diverts a fraction of capital, as in our diversion technology. But their analysis is partial equilibrium with no equilibrium default; here a positive mass of firms defaults each period, in general equilibrium.

Backup – misallocation: Hsieh–Klenow and us

	Hsieh–Klenow (2009)	This paper
Question	How large are the losses, in levels? (China and India vs the US)	How much does misallocation <i>move</i> over the business cycle?
Wedges	Two wedges per plant, τ_Y and τ_K , backed out as residuals from first-order conditions	Adverse selection $\Rightarrow$ rationing; moral hazard $\Rightarrow$ strategic default
What must be measured	Joint distribution of physical and revenue productivity, plant by plant	Two cutoffs and one loan rate
Frictionless limit	Marginal products equalized; every plant still produces	All credit to the most productive firm: $\Phi_t \rightarrow 1$, the model collapses to RBC
Evidence	Plant censuses; gains from a hypothetical reallocation	Cyclical co-movement of the 90–10 gap with growth and delinquencies

Common ground. Measured TFP splits into average firm productivity and allocative efficiency (Hsieh–Klenow, 2010); Φ_t is that second term, as an equilibrium object. Also Restuccia–Rogerson (2008), Midrigan–Xu (2014), Moll (2014), Gopinath et al. (2017), David–Venkateswaran (2019).

Backup – Hsieh–Klenow: the modeling approach (5/6)

	Hsieh–Klenow (2009)	This paper
Market structure	Monopolistic competition, CES varieties ($\sigma = 3$), constant markup	Perfect competition, one homogeneous good, constant returns, no markups
Margin of misallocation	<i>Intensive</i> : every plant operates, at a distorted size and K/L	<i>Extensive</i> : producers hold identical k and h ; cutoffs decide <i>who</i> produces
Where credit enters	As a label on τ_K : “high for firms without access to credit”	Derived: adverse selection rations credit, moral hazard prices in default
Closure	Static: aggregate K, L and the number of firms held fixed	Dynamic GE: capital accumulation, labor supply, three shocks
Efficient benchmark	Equalize marginal products; $TFP = (\sum_i A_i^{\sigma-1})^{1/(\sigma-1)}$	All credit to the most productive firm; $\Phi_t \rightarrow 1$ and the model is RBC

The sharpest line. Their τ_K is a residual that *names* the credit friction; ours is an equilibrium object – two cutoffs and a loan rate – so it responds to shocks and moves TFP period by period.

Backup – related literature (6/6)

Endogenous growth and innovation-driven reallocation.

- Endogenous growth: Lucas (1988); Romer (1990); Aghion–Howitt (1992).
- Innovation and reallocation: Akcigit–Celik–Greenwood (2016); Acemoglu–Akcigit–Alp–Bloom–Kerr (2018); Peters (2020).

How our endogenous TFP component differs.

- *Transitory*, not a permanent growth feature.
- Arises from *financial frictions*, not technology or market power.
- Reduces to two cutoffs and a price – no firm distribution to track.
- Leaves a testable cross-sectional footprint in within-industry dispersion – sidestepping a classic identification problem: aggregate time series cannot easily distinguish neoclassical from endogenous growth models (Pack, 1994), and low-frequency models are nearly indistinguishable at typical sample lengths (Müller–Watson, 2008).

Backup – full model setup

Households (KPR). Choose C , H , equity A , government bonds; supply the stochastic discount factor.

Goods firms. Two-period overlapping structure. Raise equity a_t (aliquot, independent of ω), draw ω , then sort:

- $\omega \geq \bar{\bar{\omega}}$: borrow & *produce*;
$$R_{t+1}(\omega) = \frac{\kappa_{t+1}}{Q_t} \omega + \frac{(1-\delta)}{Q_t} Q_{t+1}.$$
- $\bar{\omega} < \omega < \bar{\bar{\omega}}$: borrow & *default*; divert fraction $\Theta_t(\omega) = \omega^\psi F_t$ to the outside option (haircut ξ).
- $\omega \leq \bar{\omega}$: *lend* at inter-firm rate ρ_t .

Capital producers. Build capital with quadratic adjustment costs (ϕ), Tobin's q price Q_t .

Government. Issues bonds bought only by firms; rebates the haircut $\Xi_t = \xi D_{t-1}$ lump-sum; balanced budget.

Frictions. (i) Adverse selection: ω private $\Rightarrow$ credit rationing (Stiglitz–Weiss). (ii) Moral hazard: low- ω firms divert $\Rightarrow$ caps borrowing of even the best firm.

Shocks. AR(1) in Z (TFP), Z_I (IST), ν_c (preference).

Backup – the loan rate and the two cutoffs

Lower cutoff $\bar{\omega}_t$ (lend vs. borrow) – at $\bar{\omega}_t$, the expected return from producing equals the outside option:

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

Upper cutoff $\bar{\bar{\omega}}_t$ (divert vs. produce) – marginal firm indifferent:

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

- Default rate $\chi_t = \frac{\mu(\bar{\bar{\omega}}_t) - \mu(\bar{\omega}_t)}{1 - \mu(\bar{\omega}_t)}$.
- Three propositions + a slope condition on Θ verify the sort (no firm wants to switch segment).
- Key asymmetry: $\bar{\bar{\omega}}_{t-1}$ enters production; $\bar{\omega}_t$ is purely forward-looking.

Backup – closed-form TFP decomposition

Integrating individual output ($y_{t+1}(\omega) = Z_{t+1} \omega \tilde{k}_t^\alpha \tilde{h}_{t+1}^{1-\alpha}$, linear in ω) over producers:

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

With $\omega \sim \text{Beta}(\eta_1, \eta_2)$, use $\omega \mu'_{\eta_1, \eta_2}(\omega) \propto$ the $\text{Beta}(\eta_1 + 1, \eta_2)$ kernel and $B(\eta_1 + 1, \eta_2)/B(\eta_1, \eta_2) = \eta_1/(\eta_1 + \eta_2) \equiv c_\alpha$:

$$\Phi_t = c_\alpha \frac{1 - \mu_{\eta_1+1, \eta_2}(\bar{\bar{\omega}}_{t-1})}{1 - \mu_{\eta_1, \eta_2}(\bar{\bar{\omega}}_{t-1})} = E[\omega \mid \omega \geq \bar{\bar{\omega}}_{t-1}].$$

As $\bar{\bar{\omega}} \rightarrow 1$, both survival functions share the same $(1 - \omega)^{\eta_2}$ tail, so their ratio $\rightarrow 1/c_\alpha$ and $\Phi_t \rightarrow 1$ (RBC). $Z = 1$ at steady state $\Rightarrow \log \text{TFP}_t = \log \Phi_t + \log Z_t$.

Backup – SMM targeted moments


Quarterly, 1987:Q1–2021:Q4; HP-filtered ($\lambda = 1600$); modified optimal weighting matrix.

- Variables: real GDP, real consumption, relative price of investment, business-loan delinquency rate.
- Targeted: 4 variances, correlations, and first-order autocorrelations (14 moments total).
- Result: signs matched on all 14; data/model 90% CIs overlap on 10 (the four variances; correlations with GDP, consumption, default; GDP and consumption autocorrelations).
- Non-overlapping: $\text{corr}(\text{GDP, inv. price})$, $\text{corr}(\text{inv. price, default})$, and the two investment-price/default autocorrelations.
- Variance decomposition: TFP shocks drive GDP; *consumption* shocks drive TFP dispersion; IST least important.

Backup – parameters

<i>Conventional</i>		
β	discount factor	0.9925
α	capital share	0.30
δ	depreciation	0.01
ν	inverse Frisch	2
<i>Steady-state targets (90–10, spread, bank share, default)</i>		
η_1, η_2	Beta shape	1.543, 2.735
ξ	outside-option haircut	0.007
θ	avg. divertible fraction	0.0003
<i>SMM</i>		
ϕ	investment adj. cost	0.239
ρ_z, σ_z	TFP shock	0.407, 0.010
ρ_I, σ_I	IST shock	0.990, 0.001
ρ_ν, σ_ν	preference shock	0.986, 0.021

Backup – output, dispersion and default together




From period 2 the decaying shock lowers the expected return on capital $\Rightarrow \bar{\omega} \uparrow$, dispersion $\downarrow$, production concentrates: output and investment rise *more* than RBC; default is countercyclical.

Backlog – the two terms in the producer's return


$$R_{t+1}(\omega) = \underbrace{\frac{\kappa_{t+1}}{Q_t} \omega}_{\text{marginal product}} + \underbrace{\frac{(1-\delta)Q_{t+1}}{Q_t}}_{\text{resale}}, \quad \kappa_{t+1} \equiv \alpha Z_{t+1} (H_{t+1}/K_t)^{1-\alpha}.$$

- **The marginal product pushes up.** Z_{t+1} is still above steady state, so κ_{t+1} is higher than normal.
- **Two forces push down.** With $\rho_z = 0.41$ the technology gain is largely gone by $t+1$; and the investment boom raises K_t , which lowers κ_{t+1} for any given Z_{t+1} .
- **The resale term decides it.** Q_t jumps and $E_t Q_{t+1}$ is below it, so capital bought now is expected to resell at a loss – “it is this second term that pushes down the average return on capital”.


Backup – TFP shock, all variables




Backup – preference shock, all variables



Backup – IST shock, all variables



Backup – full TFP decomposition (3 shocks × 2 models)



Backup – with one cutoff, the endogenous component oscillates

One cutoff has to do two jobs: govern production with a one-period lag, and respond to news today.

- *Period t .* A positive shock raises the expected return to producing, so $\bar{\omega}_t$ falls, admitting less productive firms.
- *Period $t + 1$.* That lower cutoff drags down average productivity among producers, so effective TFP falls; consumption falls; R^B rises; $\bar{\omega}_{t+1}$ goes up.
- *Period $t + 2$.* The higher cutoff raises TFP; R^B falls; the cutoff falls again . . . alternating in sign and decaying at about 0.69 per period.

The root cause is a sign mismatch: the cutoff's forward-looking response and the lagged selection effect it triggers pull in opposite directions.

Backup – why the no-default model oscillates

One cutoff $\bar{\omega}_t$ must (i) govern production with a one-period lag and (ii) respond forward-looking to shocks.

- t : a positive shock raises the *expected* return to producing $\Rightarrow \bar{\omega}_t$ *falls*, admitting lower- ω producers.
- $t+1$: the lower cutoff lowers $E[\omega \mid \omega \geq \bar{\omega}_t] \Rightarrow$ effective TFP *falls*; R^B *rises*; $\bar{\omega}_{t+1}$ *up*.
- $t+2$: higher cutoff $\Rightarrow$ TFP *up*; R^B *down*; $\bar{\omega}_{t+2}$ *down* again . . . sign-alternating, decaying $\approx 0.69/\text{period}$.

Baseline fix: the defaulting mass adjusts almost immediately and diversion-correction terms add positive feedback, so the dominant eigenvalue for $\bar{\omega}_t$ stays positive – the oscillatory mode is suppressed.

Backup – variance decomposition of the Solow residual

HP-filtered ($\lambda = 1600$), 50 draws of 500 quarters, all three shocks.

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

The gap equals zero when Φ_t is constant (RBC); it measures the endogenous share, including its covariance with Z_t .

Backup – dispersion regressions (90% intervals)

	(1) 90–10 disp.		(2) 90–10 disp.	
	Data	Model	Data	Model
GDPpc growth	−2.86 [−4.16, −1.56]	−0.70 [−1.20, −0.34]	−2.72 [−3.87, −1.56]	−0.77 [−1.33, −0.33]
Lagged GDPpc growth			−2.31 [−3.56, −1.06]	−0.66 [−1.06, −0.25]
Delinquency rate	−4.22 [−5.56, −2.88]	−4.28 [−10.47, 0.04]	−4.38 [−5.57, −3.19]	−4.22 [−10.59, 0.09]
<i>N</i> (model)		34		33

Annual U.S. data 1987–2021; model = average slope over 1,000 simulated samples, 5th/95th percentiles in brackets.

Backup – delinquency regressions (90% intervals)

	(3) Delinquency		(4) Delinquency	
	Data	Model	Data	Model
GDPpc growth	−0.22 [−0.31, −0.12]	−0.11 [−0.19, −0.04]	−0.22 [−0.31, −0.13]	0.00 [−0.03, 0.03]
Lagged delinquency	0.84 [0.74, 0.94]	0.41 [0.07, 0.72]	0.78 [0.68, 0.88]	0.21 [0.02, 0.44]
Real interest rate			0.14 [0.04, 0.24]	0.26 [0.19, 0.34]
<i>N</i> (model)		34		34

Annual U.S. data 1987–2021; model = average slope over 1,000 simulated samples, 5th/95th percentiles in brackets. In spec. (4) the model's growth coefficient is essentially zero: once the real rate enters, it absorbs the cyclical variation in delinquencies that growth carries on its own in spec. (3).

Backup – why labor is hired one period ahead

The assumption. Capital *and* labor are committed at t , before ω is realized: capital because credit is rationed, labor because of “the time required to contract, train, and match workers to firms.” **What it**

buys. All producing firms are ex-ante identical, so $h_{t+1}(\omega) = \tilde{h}_{t+1}$ and $k_t(\omega) = \tilde{k}_t$. Individual output is then *linear* in ω , and aggregate output is proportional to $E[\omega \mid \omega \geq \bar{\omega}_{t-1}]$ – one number instead of a distribution. **The alternative.** If labor were hired after ω is observed, the labor–capital ratio would rise

with ω , output would be a non-linear function of ω , and the closed-form aggregation would be lost.

What would change. A within-period intensive margin: high- ω firms would expand labor on impact.

The selection margin would then share the work with a reallocation margin, and aggregation would require tracking more than the cutoff.

Backup – why a two-period overlapping structure

The device. Goods firms live two periods: they raise equity at t and produce at $t + 1$. This is familiar from decentralizations of the RBC model in which equity contracts channel household savings to firms.

What it buys. A clean separation between the moment savings are allocated (before ω is known,

hence the aliquot share a_t) and the moment the productivity draw is revealed and the firm chooses its technology. Without it, the informational friction would have to be modelled inside a long-lived firm's dynamic problem. **What it costs.** No firm-level state carries over: no net worth accumulation, no

persistence in who produces, and no entry or exit margin. Self-financing out of the friction, the mechanism at the center of much of the misallocation literature, is absent by construction – which is why the endogenous TFP here is a selection effect rather than a wealth effect.

Backup – why the Beta distribution

What is needed. Any μ on $[0, 1]$ with $\mu(0) = 0$, $\mu(1) = 1$ and $\mu' > 0$ would do for the theory. Evaluating $E[\omega \mid \omega \geq \bar{\omega}]$ in closed form requires a distributional assumption. **Why Beta.** Multiplying the

Beta(η_1, η_2) density by ω gives the Beta($\eta_1 + 1, \eta_2$) kernel, so the conditional mean collapses to a ratio of survival functions and a constant:

$$\Phi_t = c_\alpha \frac{1 - \mu_{\eta_1+1, \eta_2}(\bar{\omega}_{t-1})}{1 - \mu_{\eta_1, \eta_2}(\bar{\omega}_{t-1})}, \quad c_\alpha = \frac{\eta_1}{\eta_1 + \eta_2} = E[\omega].$$

Two shape parameters also give enough freedom to hit the dispersion and default-rate targets jointly, and bounded support matches the normalization $\omega \in [0, 1]$. **What we do not claim.** The Beta is chosen for tractability, not fitted to a firm-level productivity distribution.

Backup – why $\Theta_t(\omega) = \omega^\psi F_t$

What the theory requires. Θ_t must be non-negative, increasing, and concave in ω . Increasing, so that more productive firms can divert more and the sorting is monotone; concave, so that the slope condition evaluated at $\bar{\omega}_t$ implies it everywhere above. **The choice.** Among the many forms consistent with those requirements, we take the power specification $\Theta_t(\omega) = \omega^\psi F_t$ with $0 < \psi < 1$. F_t **is not free**. It is determined in equilibrium by the normalization that the average diverted fraction equals θ , so F_t moves with the cutoffs. A firm with a given ω can therefore divert a different share when the defaulting interval widens. **Where it comes from.** A decision problem with concealment services delivers exactly this functional form.

Backup – microfoundation of the diversion technology

A diverting firm shifts borrowed funds into private projects whose value depends on its managerial ability, itself tied to ω . Concealing them takes services from a sector whose supply is fixed in the short run.

$$\max_{\Theta} \omega^{\eta} (\Theta b)^{\gamma} a^{1-\gamma} - q_{\Theta} \Theta b \quad \Rightarrow \quad \Theta^*(\omega) = \frac{a}{b} \left(\frac{\gamma \omega^{\eta}}{q_{\Theta}} \right)^{\frac{1}{1-\gamma}} \propto \omega^{\psi}$$

Increasing and concave

With $\psi = \eta/(1-\gamma) < 1$, this is exactly the form the baseline assumes.

A single crossing


Producing pays $A_P \omega - \rho b$, linear in ω ; diverting is concave, so the two cross once.

Harder when more try it

More diverters bid up q_{Θ} against fixed supply, lowering Θ^* – the baseline's F_t .

Here the *low- ω* firms divert, around a single threshold: this justifies the shape of Θ_t , not the three-segment sort.

Backup – why bonds are sold only to firms



Letting households hold the bond would add a margin between household saving and firm equity, without changing how firms sort.

What it buys

The outside option needs an asset a diverter can hold and a creditor cannot reach.


Pricing is unaffected

The household Euler equation still pins down R_t^B : households own every asset in the end and must be indifferent at the margin.

Backup – the haircut ξ

A firm that takes the outside option earns $R_t^B - \xi$ instead of R_t^B .

$R^A = R^B$ at the steady state: households' return = risk-free rate



The diagram consists of two horizontal lines. The top line is blue and represents the risk-free rate $R^A = R^B$. The bottom line is orange and represents the return to the diverter $R_t^B - \xi$. A vertical double-headed arrow between the two lines is labeled with $\xi = 0.007$ and 'prime rate – 1-year Treasury'.

$R_t^B - \xi$ what the diverter actually earns

$\xi = 0.007$

prime rate – 1-year Treasury

Transferred, not burned

The haircut is rebated lump-sum,
 $\Xi_t = \xi D_{t-1}$: diversion moves resources
around but destroys none.

Why the prime rate

It is unsecured and defaultable – the
right counterpart to the return
households earn on equity.

The no-default variant keeps ξ at the same value, there as a pure intermediation cost in the lending spread, so the two calibrations stay comparable.

Backup – households, capital producers, government

Households. Maximize

$$E_t \sum_{\tau=0}^{\infty} \beta^{\tau} \left[\ln(C_{t+\tau} - v_{ct+\tau}) - \frac{\theta}{1+\theta} H_{t+\tau}^{1+\theta} \right] \text{ s.t. } C_t + A_t + B_t^H = R_t^A A_{t-1} + W_t H_t + R_{t-1}^B B_{t-1}^H + \Pi_t + T_t + \Xi_t.$$

Capital producers. Investment supply and capital accumulation:

$$I_t^n = Z_{It} \left[1 - \frac{\phi}{2} \left(\frac{I_t^g}{I_{t-1}^g} - 1 \right)^2 \right] I_t^g, \quad K_t = I_t^n + (1 - \delta) K_{t-1},$$

with the first-order condition delivering Tobin's q price Q_t . **Government.** Balanced budget, bonds held only by firms, haircut rebated:

$$T_t = B_t^G - R_{t-1}^B B_{t-1}^G, \quad B_t^G = B_t^H + D_t, \quad B_t^H = 0, \quad \Xi_t = \xi D_{t-1}.$$

Backup – the household problem and the preference shock

Problem. One representative household, King–Plosser–Rebelo preferences:

$$\begin{aligned} & \max_{\{A_{t+\tau}, C_{t+\tau}, H_{t+\tau}, B_{t+\tau}^H\}_{\tau=0}^{\infty}} E_t \sum_{\tau=0}^{\infty} \beta^{\tau} \left[\ln(C_{t+\tau} - \nu_{ct+\tau}) - \frac{\vartheta}{1+\vartheta} H_{t+\tau}^{1+\vartheta} \right], \\ \text{s.t. } & C_t + A_t + B_t^H = R_t^A A_{t-1} + W_t H_t + R_{t-1}^B B_{t-1}^H + \Pi_t + T_t + \Xi_t. \end{aligned}$$

The shock is a shifter inside the log: $\nu_{ct} = \rho_{\nu} \nu_{ct-1} + \varepsilon_{\nu t}$, with $\rho_{\nu} = 0.986$, $\sigma_{\nu} = 0.021$ (SMM).

First-order conditions. With λ_{ct} the multiplier on the budget constraint,

$$\lambda_{ct} = \frac{1}{C_t - \nu_{ct}}, \quad W_t = \vartheta (C_t - \nu_{ct}) H_t^{\vartheta}, \quad \lambda_{ct} = \beta E_t [\lambda_{ct+1} R_{t+1}^A] = \beta E_t [\lambda_{ct+1} R_{t+1}^B].$$

How it reaches the firms – never through technology, since Z_t is unchanged throughout:

$$\nu_{ct} \rightarrow \lambda_{ct} \rightarrow \text{SDF } \beta \frac{\lambda_{ct+1}}{\lambda_{ct}} \text{ used in every firm's problem} \rightarrow \bar{\omega}_t, \bar{\bar{\omega}}_t \rightarrow \Phi_t \rightarrow \text{measured TFP}$$

Backup – proof sketch, Proposition 1

Claim. With Θ_t non-negative and increasing, no firm with $\omega < \bar{\omega}_t$ wants to deviate from lending.

- **Lending vs. diverting.** In the loan-rate condition the left side – the payoff from lending – does not depend on the firm's own ω , while the right side – the payoff from borrowing and diverting – increases in ω through $\Theta_t(\omega)$. The two sides are equal at $\bar{\omega}_t$ by construction, so below $\bar{\omega}_t$ lending strictly dominates.
- **Lending vs. producing.** At $\bar{\omega}_t$ the expected return from producing equals the outside option $R_t^B - \xi$. Since $R_{t+1}(\omega)$ is increasing in ω and $\Theta_t(\bar{\omega}_t)b_t \geq 0$, any firm below $\bar{\omega}_t$ earns less by producing than by lending.

Ties exactly at the cutoff are resolved by assumption and have zero probability, since μ is continuous.

Backup – Proposition 1: the least productive firms lend

$\bar{\omega}_t$ ρ_t $\bar{\bar{\omega}}_t$ Φ_t

Recall that $\Theta_t(\omega)$ is the share of borrowed funds a defaulting firm can divert: non-negative and increasing in ω , so a more productive firm diverts more.

Proposition 1

If $\Theta_t(\omega)$ is non-negative and increasing in ω , firms with ω below the cutoff $\bar{\omega}_t$ have no incentive to deviate from lending.

- **Intuition.** What a firm earns by lending does not depend on its own ω – it is a pooled return – while what it would earn by borrowing and diverting *does* rise with ω .
- So if the firm at $\bar{\omega}_t$ is exactly indifferent, every firm below it strictly prefers to lend.
- Producing is worse still: their return $R_{t+1}(\omega)$ falls short of the outside option.

Backup – Proposition 2: the two cutoffs are ordered

$\bar{\omega}_t$ ρ_t $\bar{\bar{\omega}}_t$ Φ_t

Proposition 2

If $\Theta_t(\omega)$ is non-negative, then $\bar{\bar{\omega}}_t \geq \bar{\omega}_t$.

- **Intuition.** Suppose the opposite: the diversion margin would sit *below* the lending margin, leaving no interval of defaulters and making the shares in the lender's return ill-defined.
- Substituting that case into the loan-rate condition contradicts the indifference that defines $\bar{\bar{\omega}}_t$, because $R_{t+1}(\cdot)$ is increasing and equals $R_t^B - \xi$ at $\bar{\omega}_t$.
- So the three segments are always well defined, with a possibly zero mass of defaulters in between.

Backup – Proposition 3: the three-way sort

$\bar{\omega}_t$ ρ_t $\bar{\bar{\omega}}_t$ Φ_t

Proposition 3

If $\Theta_t(\omega)$ is non-negative, increasing and concave in ω , the two cutoff conditions, the loan-rate condition and the slope condition are sufficient for firms to sort into the three segments as posited.

- **Compare slopes, not levels.** At $\bar{\bar{\omega}}_t$ the payoffs are equal by construction; above it, the ranking turns on which grows faster in ω .
- Producing grows at a *constant* rate (R_{t+1} is linear in ω); diverting grows at a *decreasing* rate (Θ_t is concave in ω , so Θ'_t falls).
- The slope condition checks the hardest case, $\bar{\omega}_t$, where Θ'_t is largest: steeper there means steeper above, so the payoffs cross once.

Backup – proof sketch, Proposition 2

$\bar{\bar{\omega}}_t \geq \bar{\omega}_t$, **by contradiction**. Suppose $\bar{\bar{\omega}}_t^* < \bar{\omega}_t$. Substituting into the loan-rate condition and using $R_{t+1}(\bar{\bar{\omega}}_t^*) < R_{t+1}(\bar{\omega}_t) = R_t^B - \xi$ produces an inequality that contradicts the indifference condition defining $\bar{\bar{\omega}}_t^*$. Hence $\bar{\bar{\omega}}_t = \bar{\bar{\omega}}_t^* \geq \bar{\omega}_t$. The degenerate case $\bar{\bar{\omega}}_t = \bar{\omega}_t$, a zero mass of defaulters, is handled explicitly.

Backup – proof sketch, Proposition 3

In four steps.

- a. The payoff from diverting rises in ω while the lending payoff does not, so every firm above $\bar{\omega}_t$ prefers borrowing to lending.
- b. Differentiating the upper-cutoff condition, the diversion payoff has a flatter slope in ω than the production payoff – this is where concavity of Θ_t and the slope condition evaluated at $\bar{\omega}_t$ do the work.
- c. Combining with Proposition 2 orders the three segments.
- d. Running the same argument down from $\bar{\bar{\omega}}_t$ shows diverting dominates producing below it.

Equality of the two payoffs at $\bar{\bar{\omega}}_t$ plus the slope ranking above it gives strict dominance of producing for every $\omega > \bar{\bar{\omega}}_t$. We verify the slope condition numerically at the calibrated parameters.

Backup – equilibrium definition

A competitive equilibrium is a sequence of 22 quantities ($C, H, Y, K, B^{tot}, I^n, I^g, \Pi, T, B^G, B^H, D, \Xi, \bar{\omega}, \bar{\bar{\omega}}, Z, Z_I, \nu_C, F, L, A, B$) and 7 prices ($\lambda_C, R^A, W, R^B, Q, R, \rho$) satisfying **29 conditions**, given initial values and the shock processes. The 29 comprise four aggregation relations – aggregate output, the average

return on capital, the aggregate equity return, and the diversion factor F_t – together with 25 further conditions: the household's four first-order conditions, capital financing and labor demand, investment supply, capital accumulation, capital-producer profits, the government budget and bond-market identities, the haircut rebate, both cutoff conditions, the loan rate, aggregate lending, loan-market clearing, aggregate diversion and borrowing, goods-market clearing, Tobin's q , and the three shock processes. The slope conditions from the propositions are part of the definition: they are what make the posited sorting an equilibrium.

Backup – the frictionless limit

Claim. With $\bar{\omega}_t = \bar{\bar{\omega}}_t = 1$ and $\xi = 0$, the baseline's equilibrium conditions reduce to those of a canonical representative-firm RBC model.

- Both Beta survival functions in Φ share the same $(1 - \omega)^{\eta_2}$ tail, so their ratio tends to $1/c_\alpha$ and $\Phi_t \rightarrow 1$.
- At $\bar{\omega} = \bar{\bar{\omega}} = 1$ the lender and defaulter masses vanish, $B_t^H = D_t = 0$, and the government block is trivially satisfied.
- Aggregate output becomes $Y_t = Z_t K_{t-1}^\alpha H_t^{1-\alpha}$, and what remains are the familiar consumption and labor-supply Euler equations, factor demands, capital accumulation, Tobin's q , market clearing, and the shock processes.

The RBC benchmark used in the impulse responses is that same model, with identical preferences, adjustment costs, and shocks.

Backup – the 90–10 gap in the model

The data measure dispersion across *operating* establishments, so the model counterpart is computed over producing firms only: the Beta density is rescaled to integrate to one above the production cutoff. The 10th percentile of that truncated distribution solves

$$\int_{\bar{\omega}}^{\omega_{10}} \frac{\mu'_{\eta_1, \eta_2}(\omega)}{1 - \mu_{\eta_1, \eta_2}(\bar{\omega})} d\omega = 0.10 \iff \omega_{10} = \mu_{\eta_1, \eta_2}^{-1}[0.10(1 - \mu_{\eta_1, \eta_2}(\bar{\omega})) + \mu_{\eta_1, \eta_2}(\bar{\omega})],$$

and ω_{90} analogously with 0.90, where μ^{-1} is the Beta quantile function. The calibration target is the ratio ω_{90}/ω_{10} , matched to 2.08. Because the cutoff moves with credit conditions, so does this ratio – that is the model's counterpart of the cyclical dispersion in the data.



Backup – which shocks drive what

Share of the variance of each variable attributable to each shock (percent):

	Output	Inv. price	Consumption	Default rate	Real rate	90–10 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
Investment	1.1	0.9	2.0	2.5	2.6	3.2

- Technology shocks drive output, as usual.
- **Consumption (preference) shocks are the main driver of productivity dispersion**, not technology shocks – the endogenous channel runs through the real rate and borrowing costs.
- Investment-specific shocks matter least for the targeted moments.

Backup – how accurate is the solution?



- One random sample, all three shocks: the second-order solution tracks the fully nonlinear perfect-foresight path, while the first-order solution visibly departs from it – clearest for capital, on the right.
- The exercise validates nonlinearity in the deterministic transition, not the risk-adjustment terms that only the second-order solution prices.

Backup – a variant without strategic default

- Take away the outside option: no firm can divert, so no firm defaults.
- Two segments instead of three – lenders and borrower-producers – separated by a *single* cutoff $\bar{\omega}_t$, pinned by the same indifference condition as in the baseline.
- The government block drops out entirely ($D_t = 0$), and ξ survives only as an intermediation cost in the lending spread, held at its baseline value for comparability.
- 22 equilibrium conditions instead of 29. Households, capital producers, and the shock processes are unchanged.

Backup – the no-default model in detail

Unchanged. The household problem, the capital producers, and all three shock processes. **Removed.**

The outside option, and with it $\Theta_t(\omega)$, F_t , the aggregate diversion D_t , the rebate Ξ_t , the upper cutoff $\bar{\omega}_t$, and the government's bonds and transfers – seven variables in all. Conditions fall from 29 to 22.

Changed. Aggregate output uses the single cutoff, $Y_t = E[\omega \mid \omega \geq \bar{\omega}_{t-1}] Z_t K_{t-1}^\alpha H_t^{1-\alpha}$; the loan rate

becomes $\rho_t = R_t^B - \xi$; and aggregate lending, borrowing, and the equity return are restated for two segments. **The cutoff condition is identical** to the baseline's lower-cutoff condition, with ξ entering

the same way, which keeps the two calibrations comparable. Sorting needs only monotonicity of $R_{t+1}(\omega)$ here, rather than the three propositions plus the slope condition.