Timing Search[†]

Hanbaek Lee[‡]

University of Cambridge

May 2026

Abstract

Why do industries run on different business-cycle frequencies? I argue that cycle timing is shaped not only by shocks, but by how industries adjust. Firms choose when to enter, while permits, construction labor, and equipment suppliers cannot move instantly. This timing friction makes industries overshoot efficient scale and generate recurring cycles whose speed reflects market structure: cycle frequency takes the closed form $\omega = \sqrt{\varphi/\mu}$ in two market-structure primitives, profit curvature φ and entry-matching friction μ. In U.S. manufacturing, more geographically concentrated industries cycle faster, and the relationship is supported by a 1947 historical-geography IV with conventional first-stage $F > 10$. The model separates policies that change cycle speed from policies that reduce cyclical losses: permitting reform and cluster policy mainly retime adjustment; demand stabilization and higher bottleneck friction in the entry pipeline reduce welfare costs.

Keywords: Timing search, cycle frequency, market structure, temporal indifference, scale externality.

JEL codes: D83, E32, L16, E22.

[†]Part of this paper builds on my dissertation at the University of Pennsylvania. I thank Andrew Abel, Régis Barnichon, Vasco M. Carvalho, Harold Cole, Alessandro Dovis, John Fernald, Jesús Fernández-Villaverde, Andrew Foerster, Jeremy Greenwood, Daniel Haanwinckel, Dirk Krueger, Andrei Levchenko, Huiyu Li, Christian Moser, Makoto Nirei, Guillermo Ordoñez, William Peterman, Nicolas Petrosky-Nadeau, Andrew Postlewaite, José Víctor Ríos Rull, Anna Sanktjohanser, and Frank Schorfheide for helpful comments and discussions, as well as seminar participants at the University of Pennsylvania, Bilkent University, the Federal Reserve Banks of Boston, San Francisco, and St. Louis, the Federal Reserve Board, the Society for Economic Dynamics, the University of Cambridge, and the University of Tokyo. All errors are my own.

[‡]Faculty of Economics, University of Cambridge, Sidgwick Avenue, Cambridge CB3 9DD, United Kingdom. Email: hl610@cam.ac.uk.

1 Introduction

Why do different industries fluctuate at different frequencies? In U.S. data, spectral peak periods range from 3.0 to 4.7 years in industrial production and from 2.4 to nearly 10 years in establishment entry, while industries share the same monetary policy, aggregate demand shocks, and financial conditions. Standard business-cycle frameworks derive cycle properties from a common shock process (Kydland and Prescott, 1982) or policy rule, leaving cross-industry frequency variation unexplained. The obstacle is that no existing structural model maps the observed cross-industry dispersion to industry primitives: shock-driven business-cycle models predict that industries hit by common shocks share a common cycle length, while endogenous-cycle models generate internal persistence without tying cycle frequency to market-structure primitives or separating the determinants of timing from those of welfare. This paper addresses the gap by introducing *timing search*: agents choose *when* to enter under matching frictions that penalize rapid adjustment. Indifference across entry dates in equilibrium leads to oscillation rather than monotone return to optimal scale, with cycle frequency taking the closed form

$$\omega \;=\; \sqrt{\varphi/\mu}, \tag{1}$$

where φ is the curvature of profits around optimal industry scale and μ is the entry-matching friction—both market-structure primitives, not properties of the shock process.

This paper delivers three key results. First, I propose a mechanism for the industry-specific cycles documented in U.S. data (Theorem 1, Section 3). The equilibrium concept underlying this result is *time-indifference equilibrium*, an intertemporal no-arbitrage condition for entry timing whose closest formal precedents are Hotelling (1931) (intertemporal arbitrage) and Burdett and Mortensen (1998) (equal-value condition across continuous choices). The novelty is that combining this indifference condition with matching frictions—which make the value of entry at date t depend on the rate of net adjustment $\dot{n}(t)$—translates a forward-looking value equation into a constraint on the equilibrium time path, and this constraint admits the closed-form harmonic solution (1). With slow-moving (φ, μ), realized cycles are irregular but their average dominant frequency is ranked by the persistent component of φ/μ (Section 7). A distinctive prediction is that the cycle period depends on market structure alone, not on cycle amplitude—so deeper recessions recover at proportionally faster speeds in the same time as shallow ones. This matches the plucking pattern of Friedman (1993), confirmed in U.S. industrial production over 1972–2024.

Second, I bring the model's prediction that cycle frequency is pinned by market structure—not the shock process—to U.S. industry data (Section 5). Across 15 manufacturing 3-digit

industries, the spectral peak frequency varies systematically with the Ellison and Glaeser (1997) geographic concentration index, an empirical proxy for profit curvature: industries with stronger agglomeration cycle at higher frequencies. A historical-geography IV built from 1947 state-level specialization net of one-county dominance has a robust first-stage statistic above 10 and yields a cross-industry slope consistent with the structural prediction; primitive Kim and County Business Patterns instruments give the same sign, and an independent BDS subsector cross-section does as well.

This evidence is consistent with a profit-curvature channel: across manufacturing, entry inputs (permits, construction crews, equipment supply chains) flow relatively freely, so the matching friction varies little across 3-digit industries, and the cross-industry signal in cycle frequency comes from differences in the profit landscape. Three additional diagnostics are consistent with the mechanism but rest on weaker identification: BDS net-entry rates show oscillatory features outside the AR(1) class in 10 of 19 sectors, manufacturing and oil investment oscillate at industry-specific frequencies after large historical shocks, and 12 of 19 BDS sectors overshoot above their pre-pandemic trend after COVID-19 by at least one pre-COVID standard deviation. Among potential alternatives—AR(1) RBC, NK AR(2), Beaudry et al. (2020), and multi-sector RBC with input-output linkages—each captures part of the joint pattern but not the cross-industry slope at observed magnitudes; reproducing the full pattern requires a mechanism that maps industry-level primitives into industry-level frequencies (Online Appendix F).

Third, I identify which entry-side policies move the welfare cost of cycles and which do not, delivering a sharp policy taxonomy (Theorem 3, Section 6.2). The taxonomy follows from a frequency-welfare separation that the model delivers: $\omega_0 = \sqrt{\varphi/\mu}$ depends only on the market-structure pair (φ, μ), while the welfare loss $\mathbb{E}[\mathcal{L}] = \sigma_\varepsilon^2/(2\chi)$ depends only on the disjoint pair of profit-shock variance and entry-pipeline friction χ. Two natural classes of policy therefore fall on the timing-only side: streamlining entry matching (e.g., faster permitting) and strengthening agglomeration (e.g., cluster subsidies) change cycle speed without affecting welfare cost. Only policies that raise the bottleneck friction in the entry pipeline (higher χ) or stabilize profit shocks (lower σ_ε) move welfare. The taxonomy thus distinguishes welfare-improving policy from policy that, while visibly altering cycle dynamics, leaves welfare unchanged. This is directly relevant to the recent industrial-policy turn—the U.S. CHIPS and Science Act, the Inflation Reduction Act, the EU Net-Zero Industry Act (Juhász et al., 2024)—where the framework provides a structural lens for asking which proposed entry-side instruments actually move welfare and which merely reshape cycle dynamics. Under the per-firm welfare metric the model adopts, many CHIPS- and IRA-style "capacity expansion" subsidies that reduce supplier-side friction can in fact *raise* welfare cost

by removing the dissipation that bounds cycle amplitude (Section 6.2).

Related literature This paper builds on four literatures: structural and endogenous cycles, search and matching frictions, industry dynamics, and agglomeration.

Structural and endogenous cycles. The structural-cycle tradition is anchored by Friedman (1993), who reads recessions as plucks below an unobserved natural ceiling rather than as symmetric stochastic deviations; the deeper-recession-faster-recovery pattern that view predicts is revisited empirically by Bordo and Haubrich (2017) and Dupraz et al. (2019), and emerges in the present model as a corollary of isochrony (Section 3.5). The closest formal macroeconomic antecedent is Beaudry et al. (2020), who show that financial frictions can generate endogenous boom-bust dynamics through a stochastic limit cycle (extended in Beaudry et al. (2024)). Beaudry et al. (2020) obtain a true limit cycle from a Hopf bifurcation, so the existence of cycles is itself a structurally stable prediction of their framework. The deterministic timing-search benchmark in this paper is a center rather than a limit cycle, so cycles are permitted but not selected by the deterministic dynamics; bottleneck damping ($\chi > 0$) and shocks ($\sigma_\varepsilon > 0$) jointly select the stationary amplitude in the stochastic equilibrium (Propositions 3–4). The present paper therefore trades the cycle-robustness of a Hopf attractor for a closed-form mapping from market-structure primitives to cycle frequency and a frequency-welfare separation theorem that Beaudry et al. (2020)'s framework does not deliver; the two approaches answer different questions through different mathematical structures. The animal-spirits and self-fulfilling-prophecies tradition of Farmer and Guo (1994) and Farmer (1999) generates endogenous fluctuations from indeterminacy and belief shocks; here, by contrast, cycle frequency is pinned by market-structure primitives, and in the stochastic equilibrium the cycle amplitude is pinned by bottleneck damping and profit shocks (stationary severity $\sigma_x = \sigma_\varepsilon/\sqrt{2\chi\varphi}$) rather than by self-fulfilling beliefs—the deterministic frictionless model is a center with free amplitude, made determinate by $\chi > 0$ and $\sigma_\varepsilon > 0$ in Theorem 2 (of which Propositions 3–4 are the two-channel preview). The frequency-vs-welfare disjointness is reserved for Theorem 3: only there does φ cancel from the welfare formula $\mathbb{E}[\mathcal{L}] = \sigma_\varepsilon^2/(2\chi)$, while it remains a determinant of amplitude. Beaudry and Portier (2006) document stock prices leading TFP; Grandmont (1985), Azariadis (1981), Benhabib and Farmer (1994), and Goodwin (1967) provide earlier endogenous-cycle precedents. The closest timing-based antecedent is Jovanovic (2009), who shows how investment timing can generate endogenous business cycles; the present paper shifts the object from the option value of investment timing to the equilibrium distribution of entry dates, which yields a closed-form mapping from market-structure primitives to cycle frequency.

Search and matching frictions. Timing search is in the spirit of Stigler (1961), McCall

(1970), and especially Burdett and Mortensen (1998): a temporal counterpart in which agents choose among dates rather than among contemporaneous alternatives, with an ironing condition analogous to the equal-value condition in Burdett and Mortensen (1998) (Online Appendix A.1). The micro-foundation $\mu = (1-\alpha)/[\alpha(\delta n_p)^2]$ further connects to the Diamond-Mortensen-Pissarides literature (Diamond, 1982; Mortensen and Pissarides, 1994; Pissarides, 2000; Petrongolo and Pissarides, 2001; Shimer, 2005): matching frictions in entry generate equilibrium business cycles in the same way DMP matching frictions generate equilibrium unemployment.

Industry dynamics. The model isolates a timing-search channel within an endogenous entry environment, complementary to the richer heterogeneity and selection mechanisms of Hopenhayn (1992) and Hopenhayn and Rogerson (1993); the entry-wave dynamics generated by the ironing condition complement Jovanovic (1982)'s learning-based exit explanation, and Davis and Haltiwanger (1992) document the entry/exit empirics this literature studies. A separate strand on aggregate fluctuations from disaggregated origins (Gabaix, 2011; Acemoglu et al., 2012) provides the natural alternative the paper tests against in Online Appendix F.2.

Agglomeration. The scale-congestion logic of Krugman (1991), Duranton and Puga (2004), Rosenthal and Strange (2004), Ellison and Glaeser (1997), and Ellison et al. (2010) is transposed from space to time: combined with timing frictions, it generates temporal patterns (cycles) rather than spatial patterns (clusters), and the EG concentration index serves as an external proxy for the sharpness of the scale-congestion trade-off in the present framework; Syverson (2004) provides direct evidence of the productivity-density relationship that underlies the hump-shaped profit assumption.

2 Environment

The environment has three elements chosen to make the cycle-frequency result precise: a hump-shaped profit landscape that gives the industry a well-defined optimal scale, a matching function for entry that endogenously generates a quadratic adjustment cost, and linear-utility households that let industries evolve independently. Section 2.1 sets up the agents, Section 2.2 the profit landscape, Section 2.3 the timing friction, and Section 2.4 combines them into the reduced-form dynamics that drive the equilibrium.

2.1 Agents

The economy consists of two types of agents: a representative household that consumes and supplies labor, and a fixed mass of potential entrepreneurs who alternate between

operating firms and searching for entry opportunities. The household has linear utility over consumption, $U = \int_0^\infty e^{-\rho t} C(t)\, \mathrm{d}t$, and supplies labor perfectly elastically at a fixed wage $\bar{w}$.[1] Three ingredients define the entry side of the environment.

Assumption 1 (Entry economy)**.**
A mass $M > 0$ of potential entrepreneurs, each at any instant either active *(operating a firm) or* searching *(a potential entrant). Active firms produce using labor hired from the representative household at wage $\bar{w}$, with technology $y_i = z l_i^\gamma$, $\gamma \in (0,1)$. Each entry incurs a per-firm sunk cost $\kappa \geq 0$, paid at the moment of entry. Firms exit at Poisson rate $\delta > 0$, returning to the search pool.*

Let $n(t)$ denote the mass of active production units at date t and $S(t) = M - n(t)$ the mass of searchers. Throughout the paper, I interpret n as the mass of active *establishments*—physical production units such as plants, stores, or branches—rather than firms (legal entities). The establishment interpretation matches both empirical series used in Sections 5–5.4: BEA investment in structures measures the rate of new establishment construction, and BDS establishment entry and exit rates measure establishment births and deaths. The matching-congestion microfoundation of μ (Section 2.3) is also most natural at the establishment level, where bottlenecks in permits, specialist labor, equipment, and sites constrain rapid adjustment. The aggregate state of the economy is fully characterized by $n(t)$ and its rate of change $\dot{n}(t)$.

2.2 Matching efficiency and the profit landscape

Each firm's effective output depends on market thickness. More firms create denser networks of suppliers, customers, and complementary services, improving efficiency. But beyond a point, additional firms congest the market. A scale efficiency function $m(n)$ multiplies each firm's output: $y_i = z l_i^\gamma \cdot m(n)$. Profit is $\pi(n) = \bar{\pi}(w) \cdot m(n)$, where $\bar{\pi}$ is the standard profit given the wage.[2]

Assumption 2 (Gaussian scale efficiency)**.**
Scale efficiency takes the Gaussian form:

$$m(n) = \exp\left(-\frac{\varphi}{2}(n - n_p)^2\right), \qquad \varphi > 0,\ n_p > 0. \tag{2}$$

[1] Equivalently, the household has linear utility in consumption and linear disutility from labor, $U = \int_0^\infty e^{-\rho t}\big[C(t) - \bar{w}\, L(t)\big]\, \mathrm{d}t$, so the static labor-supply condition pins $w = \bar{w}$ and labor is supplied elastically. This is in the spirit of Greenwood et al. (1988) preferences in eliminating the wealth effect on labor supply.

[2] Solving the firm's static problem $\max_l z l^\gamma - w l$ gives optimal labor $l^* = (\gamma z / w)^{1/(1-\gamma)}$ and standard profit $\bar{\pi}(w) = (1-\gamma)\gamma^{\gamma/(1-\gamma)} z^{1/(1-\gamma)} w^{-\gamma/(1-\gamma)} > 0$. With $m(n_p) = 1$ in the Gaussian specification (Assumption 2), the constant π_0 in equation (3) equals $\bar{\pi}(\bar{w})$.

The function peaks at $n = n_p$ (optimal market scale) and falls off symmetrically. The parameter φ governs the sharpness of the trade-off: high φ means a narrow peak (sharp trade-off); low φ means a broad peak (robust to deviations). The hump-shaped specification is consistent with the productivity-density relationship documented in Syverson (2004), in which establishment-level productivity rises and then falls with market-level density—the empirical pattern this paper's profit landscape encodes. Online Appendix A.2 provides a matching-theoretic microfoundation.

For the analytical results, I work with the reduced-form Gaussian profit:

$$\pi(n) = \pi_0 \exp\Big(-\frac{\varphi}{2}(n - n_p)^2\Big), \qquad \pi_0 > 0. \tag{3}$$

The semi-elasticity is exactly linear:

$$f(n) \equiv \frac{\pi'(n)}{\pi(n)} = -\varphi(n - n_p). \tag{4}$$

This linearity delivers exact harmonic oscillations (Theorem 1). The general case—where wage congestion adds a non-quadratic correction to $\ln \pi$—produces non-harmonic oscillations with the same leading-order frequency (Online Appendix B.3).

2.3 The timing friction: congestion in the entry process

Entry is mediated by a Cobb-Douglas matching technology $m(S, R) = S^\alpha R^{1-\alpha}$ with elasticity $\alpha \in (0, 1)$, where the $S(t) = M - n(t)$ searchers compete for $R(t)$ units of *entry resources*—permits, specialized labor, venture capital, factory sites—to form matches. The resource stock R is endogenously supplied by an infrastructure sector; the matching-friction derivation in this subsection uses R at its steady-state level R_{ss} to pin the timing friction μ. Deviations of R from R_{ss}—driven by net adjustment pressure $\dot{n}$ in the stationary cycle—generate an additional dissipation channel that pins the cycle amplitude, taken up in Section 4.1. The matching infrastructure is tuned to clear the steady-state throughput $e_{ss} = \delta n_p$, with routine replacement at rate $\delta n(t)$ serviced smoothly. The timing friction operates on net adjustment pressure $\dot{n}$—deviations of the adjustment-speed margin from zero: rapid entry ($\dot{n} > 0$) overloads permitting and construction, while rapid net contraction ($\dot{n} < 0$) overloads resale and reorganization. Both reduce the matching probability per searcher and produce the timing friction.

Assumption 3 (Timing friction)**.**

The multiplicative timing friction takes the log-quadratic form

$$\xi(\dot{n}) = e^{-\frac{\mu}{2}\dot{n}^2}, \qquad \mu > 0,$$

where μ is the friction parameter.

The functional form arises as the continuous-time limit of *proportional adjustment costs* in the entry pipeline: in a discrete-time economy where effective output is reduced by a fraction proportional to the squared adjustment rate $(\Delta n/\Delta)^2$ each period, taking the period length $\Delta \to 0$ yields the exponential-quadratic form, with μ a free friction parameter (Online Appendix A.3). The exponential-quadratic specification is canonical: any symmetric friction ξ with $\xi(0) = 1$ and $\xi'(0) = 0$ admits the second-order expansion $\ln \xi(\dot{n}) = -\frac{\mu}{2}\dot{n}^2 + O(\dot{n}^4)$. Symmetry of the flow penalty in $\dot{n}$ is itself a knife-edge condition: under asymmetric strain, the leading-order ironing condition would carry a cubic term in $\dot{n}$ that produces amplitude-dependent frequency (Online Appendix B.5); the exact-harmonic result of Theorem 1 requires the symmetric specification, while the qualitative cycle survives asymmetric strain.

Combining Assumption 3 with the Cobb-Douglas matching technology described above pins the friction parameter to a specific function of structural primitives: any symmetric strain term $g(\dot{n}^2/(\delta n_p)^2)$ with $g(0) = 0$ and $g'(0) = 1/2$ multiplying the matching probability gives, at second order around $\dot{n} = 0$,

$$\mu = \frac{1-\alpha}{\alpha\,(\delta n_p)^2}, \tag{5}$$

independent of the higher-order shape of g (Online Appendix A.4). Here α is the matching elasticity and δn_p is the steady-state replacement throughput. Equation (5) is the calibration I adopt when reading μ off the matching microfoundation; the body's results that depend on μ alone (e.g., $\omega = \sqrt{\varphi/\mu}$) hold under Assumption 3 for any positive friction parameter.

The timing friction μ decreases in the matching elasticity α and decreases in the steady-state throughput δn_p (the entry-resource stock R is absorbed into the steady-state calibration of n_p via the matching technology and therefore does not appear in (5) as a separate factor). Effective output is $y(t) = \pi(n(t))\,\xi(\dot{n}(t))$. The quadratic form in $\dot{n}$ arises generically from any smooth symmetric matching friction; Online Appendix A.3 derives the same ξ from discrete-time proportional adjustment costs.

2.4 Combined structure

Combining the Gaussian profit and the Gaussian friction, effective output is:

$$y(t) = \pi_0 \exp\Big(-\frac{\varphi}{2}(n - n_p)^2 - \frac{\mu}{2}\dot{n}^2\Big). \tag{6}$$

The effective output is a bivariate Gaussian kernel in $(n, \dot{n})$, centered at $(n_p, 0)$. Both terms impose quadratic penalization of deviation from the optimum: n_p for market scale and $\dot{n} = 0$ for adjustment speed.

3 Ironing along time: equilibrium and the main result

The equilibrium condition that agents cannot profit by retiming their entry—the ironing condition—pins a rate of entry along any approach to optimal scale. That pinned rate becomes momentum the industry cannot discard at the optimum without creating arbitrage. The result is cycles, and in the frictionless limit they are exactly harmonic with frequency $\omega = \sqrt{\varphi/\mu}$. This section formalizes the argument.

3.1 The agent's problem

A searching agent chooses a date τ at which to enter. Upon entry, she operates a firm earning effective flow payoff $\pi(n(t))\xi(\dot{n}(t))$ per unit time, discounted at rate $r \geq 0$, with exogenous exit at rate $\delta > 0$. The value of entering at date τ is:

$$V(\tau) = \int_{\tau}^{\infty} e^{-(r+\delta)(t-\tau)}\, \pi(n(t))\, \xi(\dot{n}(t))\, \mathrm{d}t. \tag{7}$$

$V(\tau)$ is the value of a single entry spell—operating from τ until stochastic exit. Re-entry after exogenous exit is treated as a fresh draw of the same problem for whichever agent occupies the search pool at that date, so the single-spell formulation is without loss for equilibrium dynamics under the ironing condition. The agent maximizes $V(\tau)$ over τ, taking the aggregate path $\{n(t)\}$ as given. An agent who does not enter at date t remains in the search pool, with continuation value

$$V^{\text{search}}(t) = \max_{\tau \geq t} e^{-r(\tau - t)}\big[V(\tau) - \kappa\big], \tag{8}$$

where κ is the per-firm sunk entry cost from Assumption 1, paid at the moment of entry.

3.2 Equilibrium

An *entry policy* is a measurable function $e : \mathbb{R}_+ \to \mathbb{R}_+$ specifying the flow rate of entrants at each date. Together with the initial condition $n(0)$, it determines the entire path of the economy through the aggregate consistency condition $\dot{n} = e - \delta n$.

Definition 1 (Timing search equilibrium)**.**
A timing search equilibrium *is a path* $\{n(t)\}_{t\geq 0}$ *and an entry policy* $e(\cdot)$ *such that:*

(i) ***Individual optimality****: Each agent chooses her entry date to maximize* $V(\tau)$*.*

(ii) ***Temporal free entry (with complementary slackness)****:* $V(t) \leq V^{search}(t) + \kappa$ *at all dates, with equality whenever* $e(t) > 0$ *and* $n(t) < M$*. At dates where the population cap binds* ($n(t) = M$) *or the entry flow is zero* ($e(t) = 0$)*, the inequality may be strict.*

(iii) ***No profitable timing deviation****:* $V(\tau) \leq V(t)$ *for all* τ *and all* t *in the support of entry.*

(iv) ***Aggregate consistency****:* $\dot{n}(t) = e(t) - \delta\, n(t)$*, where* $e(t) \geq 0$*.*

(v) ***Feasibility****:* $n(t) \leq M$ *for all* t*.*

The no-profitable-timing-deviation condition Condition (iii) is the key requirement. In Burdett and Mortensen (1998), the analogous condition is that no firm can profit by reposting its wage, which yields equal profits across wages. Here, the condition that no agent can profit by retiming her entry yields equal profits across dates. If any date offered strictly higher value, agents would crowd that date, changing $n(t)$ until the advantage was eliminated.

Commitment versus dynamic optimization The formulation in (7)–(8) is a commitment problem: agents choose τ taking the entire future path $\{n(t)\}$ as given. An alternative formulation would have agents observe the current state $(n(t), \dot{n}(t))$ at each instant and decide dynamically whether to enter. The ironing condition makes the two equivalent. Since $V(\tau) = V^*$ for all τ in equilibrium, an agent who could re-optimize at any date t would find no profitable deviation—she is indifferent across all entry dates. The commitment equilibrium is therefore also an equilibrium of the dynamic problem: agents who can freely re-time their entry choose not to.

Equation of motion The equilibrium dynamics of $(n_t, \dot{n}_t)$ are governed by a single equation of motion that applies throughout the paper:

$$\ddot{n} \;+\; \frac{\chi}{\mu}\,\dot{n} \;+\; \frac{\varphi}{\mu}\,(n - n_p) \;=\; \frac{\sigma_\varepsilon}{\mu}\,\varepsilon_t, \tag{9}$$

where $\chi \geq 0$ is the bottleneck friction (introduced in Section 4.1), $\sigma_\varepsilon \geq 0$ is the profit-shock intensity (introduced in Section 4.3), and ε_t is Gaussian white noise. The remainder of this section works with the deterministic frictionless specialization $\chi = \sigma_\varepsilon = 0$, under which the ironing condition (Proposition 1) reduces (9) to the harmonic oscillator (13) of Theorem 1. Section 4 restores positive χ and σ_ε, under which (9) becomes the damped stochastic oscillator (19) of Proposition 4. The unified statement permits a single equilibrium concept (Definition 2 below) that applies in both regimes.

Definition 2 (Markov stationary timing-search equilibrium)**.**
The deterministic equilibrium of Definition 1 extends to the stochastic environment of Section 4 as follows. A Markov stationary timing-search equilibrium *is a triple* (μ^*, e^*, V^*)*—a stationary distribution* μ^* *of* $(n, \dot{n})$ *with finite second moments, an entry policy* e^**, and a value function* V^**—satisfying:*

(i) ***Aggregate consistency****: under* e^**, the state process generated by (9) admits* μ^* *as its invariant law, and* $\dot{n}_t = e^*_t - \delta n_t$ *along paths.*

(ii) ***Stochastic ironing (martingale form)****: at every* $(n, \dot{n}) \in \operatorname{supp} \mu^*$*, the value function* V^* *satisfies the continuous-time Bellman equation*

$$\mathbb{E}\big[dV^*(n_t, \dot{n}_t) \,\big|\, \mathcal{F}_t\big] \;=\; \big[(r+\delta)\, V^*(n_t, \dot{n}_t) \;-\; \pi(n_t)\, \xi(\dot{n}_t)\big]\, dt. \tag{10}$$

Equivalently, the discounted-value-plus-flow process $e^{-(r+\delta)t}\, V^*(n_t, \dot{n}_t) + \int_0^t e^{-(r+\delta)s}\, \pi(n_s)\, \xi(\dot{n}_s)\, ds$ *is a martingale under* μ^**. In the* $\sigma_\varepsilon \downarrow 0$ *limit, (10) reduces pathwise to the deterministic ironing condition of Proposition 1.*

(iii) ***No-deviation across the stationary support****: for every two stopping times* τ_1, τ_2 *supported on* μ^**, the expected option value is equalized,*

$$\mathbb{E}\big[V^*(n_{\tau_1}, \dot{n}_{\tau_1})\big] \;=\; \mathbb{E}\big[V^*(n_{\tau_2}, \dot{n}_{\tau_2})\big] \;=\; \mathbb{E}_{\mu^*}[V^*];$$

no agent profits in expectation by retiming entry from one stopping time to another, the stochastic counterpart of the deterministic indifference $V(\tau) \equiv V^*$ *along a single path.*

Clause (iii) is the substantive new restriction: among the infinite family of stationary distributions admissible under (i)–(ii), it selects those at which conditional-expected option value is the same across the states the economy visits in the long run. The selection result of Section 4 (Theorem 2) shows that exactly one such distribution exists, and that the deterministic equilibrium of Theorem 1 is its vanishing-noise limit. Note that under the linear-Gaussian

dynamics of (9), V^* is a non-constant quadratic function of $(n - n_p, \dot{n})$, and the equilibrium value level $\mathbb{E}_{\mu^*}[V^*]$—not V^* itself—is the conserved object across stationary entry dates.

3.3 The ironing condition

The key equilibrium condition is that agents cannot profit by retiming their entry. This is the temporal analog of the equal-profit condition in Burdett and Mortensen (1998): Burdett-Mortensen equalize profits across wages so that no firm gains by reposting, and timing search equalizes payoffs across dates so that no agent gains by retiming. I call this the *ironing condition*, since it "irons out" any temporal variation in the value of entry.

Proposition 1 (Ironing condition)**.**
In a timing search equilibrium where entry occurs at all dates, agents are indifferent across entry dates. Equivalently, the effective flow payoff is equalized:

$$\pi(n(t))\,\xi(\dot{n}(t)) = C \qquad \text{for all } t, \tag{11}$$

for some constant $C > 0$.

Proof. Suppose entry occurs at all dates. For any two dates τ_1, τ_2 in the support of entry, condition (iii) of Definition 1 requires $V(\tau_1) \leq V(\tau_2)$ and $V(\tau_2) \leq V(\tau_1)$, so $V(\tau_1) = V(\tau_2)$. Since the support is all of $\mathbb{R}_+$, $V(\tau) = V^*$ for all $\tau \geq 0$. The equilibrium path $n(\cdot)$ is C^1 by aggregate consistency (condition (iv)) and the smoothness of π and ξ, so $V(\tau)$ is differentiable in τ.[3] Differentiating (7):

$$\frac{\mathrm{d}V}{\mathrm{d}\tau} = -\pi(n(\tau))\,\xi(\dot{n}(\tau)) + (r+\delta)\,V(\tau).$$

Setting $\frac{\mathrm{d}V}{\mathrm{d}\tau} = 0$ (since V is constant): $\pi(n(\tau))\xi(\dot{n}(\tau)) = (r+\delta)V^* \equiv C$. ■

Regularity condition The condition that entry occurs at all dates is satisfied whenever the searcher pool $S(t) = M - n(t)$ is bounded away from zero along the equilibrium path and the timing friction μ is positive. With a large population M and moderate cycle amplitude $A < M - n_p$, the condition holds. Online Appendix B.4 provides a formal verification.

[3] The differentiability of V in τ requires that $t \mapsto \pi(n(t))\xi(\dot{n}(t))$ is locally integrable along the equilibrium path, which holds whenever $n(\cdot) \in C^1$. This regularity is guaranteed by the smoothness of π (Assumption 2) and ξ (Section 2.3), combined with the continuity of the entry flow $e(\cdot)$ in condition (iv).

3.4 The main result

The equilibrium under Gaussian profits and zero bottleneck friction is an exact harmonic oscillator with frequency $\omega = \sqrt{\varphi/\mu}$ (Theorem 1). This subsection derives that result. The frictionless case is not the most realistic specification, but it is the most analytically tractable one: the ironing condition reduces to a frictionless harmonic oscillator whose orbit and period admit closed-form expressions at every amplitude.[4] Taking logarithms of the ironing condition (11), substituting the functional forms (6), and rearranging:

$$\underbrace{\frac{\varphi}{2}(n - n_p)^2}_{\text{profit loss}} + \underbrace{\frac{\mu}{2}\dot{n}^2}_{\text{adjustment cost}} = \underbrace{E}_{\substack{\text{equilibrium entry} \\ \text{value (constant)}}}, \tag{12}$$

where $E = \ln \pi_0 - \ln C > 0$ is the equilibrium value level pinned by the ironing condition (equal value across all entry dates). Differentiating Equation (12) with respect to time gives $\dot{n}\,[\varphi(n - n_p) + \mu\ddot{n}] = 0$; dividing through by $\dot{n}$, which is nonzero along the orbit, yields the corresponding law of motion:

$$\underbrace{\ddot{n}}_{\substack{\text{change in} \\ \text{entry rate}}} = \underbrace{-\frac{\varphi}{\mu}\,(n - n_p)}_{\substack{\text{pull toward} \\ \text{optimal scale}}}. \tag{13}$$

At every instant the change in the entry rate is pinned by the current deviation from optimal scale, scaled by the ratio of profit curvature to timing friction φ/μ: industries far from n_p accelerate quickly back toward it, industries near n_p adjust slowly. This is the deterministic skeleton of the timing-search dynamics. Section 4 adds congestion damping and productivity shocks to the same skeleton, but the restoring-force coefficient φ/μ that pins the natural frequency in Theorem 1 is unchanged through both extensions.

Theorem 1 (Exact harmonic oscillations)**.**
In a timing search equilibrium (Definition 1) under Assumptions 1–3, and provided that entry occurs at all dates (the regularity condition formalized in the verification paragraph below, equivalent to $A \leq \delta n_p/\sqrt{\omega^2 + \delta^2}$), the equilibrium path is:

$$n(t) = n_p + A\cos(\omega t + \phi), \qquad \omega = \sqrt{\frac{\varphi}{\mu}}. \tag{14}$$

[4] Section 4 extends this frictionless oscillator with bottleneck friction and stochastic productivity shocks; the same closed-form natural frequency $\omega = \sqrt{\varphi/\mu}$ carries over to the stationary cycle of that section, where it locates the spectral peak of $n(t)$ up to a small bottleneck-damping correction (Proposition 4).

The solution is exact within the feasible all-entry region and requires no linearization. The amplitude A and phase ϕ are determined by initial conditions (subject to the feasibility bound), and the ironing condition holds with $E = \frac{\varphi}{2}A^2$.

Proof. The ironing condition (Proposition 1) gives (12). Differentiation yields the harmonic oscillator equation (13), whose general solution is (14) with $\omega = \sqrt{\varphi/\mu}$. The solution is exact because the Gaussian specification makes the semi-elasticity $f(n)$ exactly linear (equation (4)), so no Taylor approximation is needed. Substituting back: $\frac{\varphi}{2}A^2\cos^2(\omega t+\phi) + \frac{\mu}{2}\omega^2 A^2 \sin^2(\omega t + \phi) = \frac{\varphi}{2}A^2 = E$. ■

Verification and scope The pair $(n(t), e(t))$ is a timing search equilibrium for any amplitude $A \leq \delta n_p/\sqrt{\omega^2+\delta^2}$, the binding sufficient condition for $e(t) = \dot{n} + \delta n \geq 0$ throughout the cycle (Online Appendix B.6). Industries with net entry routinely $|\dot{n}| > \delta n_p/2$ are outside the regime where the ironing condition holds globally. The center structure leaves A free; the stochastic extension and dissipation select the unique stationary equilibrium of Theorem 2 from primitives disjoint from those pinning the period.

Structural decomposition The frequency $\omega = \sqrt{\varphi/\mu}$ decomposes into two primitives. The numerator φ is the curvature of the profit landscape: how sharply profits fall when the industry deviates from optimal scale. The denominator μ is the timing friction: how costly it is to adjust the number of active firms rapidly, arising from congestion in the entry matching process. Treated as an abstract pair, (φ, μ) alone determines ω; the frequency does not depend on productivity z (which affects profit levels but not curvature), on the realized exit rate δ, or on the shock process (Section 4). When μ is calibrated structurally via Assumption 3, $\mu = (1-\alpha)/[\alpha(\delta n_p)^2]$ inherits dependence on the matching elasticity α and the steady-state throughput δn_p; the resource-stock R underlying the matching technology is absorbed into the steady-state calibration of n_p and does not appear separately in the formula. The structural decomposition is therefore: ω depends on (φ, μ) as primitives, and on (α, δ, n_p) when μ is read off the microfoundation. Substituting $\mu = (1-\alpha)/[\alpha(\delta n_p)^2]$ into $\omega = \sqrt{\varphi/\mu}$ gives the structural frequency

$$\omega = \delta n_p \sqrt{\alpha\varphi/(1-\alpha)}, \tag{15}$$

connecting cycle frequency to three independently measurable objects: matching elasticity α, profit curvature φ, and steady-state turnover δn_p. In the absence of direct estimates of entry-matching elasticities, I use the labor-market range $[0.5, 0.7]$ from Petrongolo and Pissarides (2001) as a benchmark; the structural frequency is robust to $\alpha \in [0.3, 0.8]$ since μ

varies smoothly with α over this range. Table I contrasts this decomposition with the way standard frameworks pin cycle frequency.

Multi-industry general equilibrium Because the wage is fixed and households have linear utility, each industry's profit landscape depends only on its own number of active firms. Under idiosyncratic shocks and no inter-industry input-output linkages or demand spillovers, a multi-industry economy with I industries each characterized by (φ_i, μ_i) evolves as I independent timing-search systems, with aggregate spectral density a superposition of industry-specific peaks at frequencies $\omega_i = \sqrt{\varphi_i/\mu_i}$.[5] The isolation of frequency from the shock process is therefore a *cross-sectional* prediction: different industries cycle at different frequencies determined by their respective market structures. Under common shocks or input-output linkages, the spectral superposition does not survive aggregation. Section 7 bounds the scope of the closed-form result, with quantitative leakage of (φ, μ) into the welfare cost under CRRA preferences documented in Online Appendix E.1.

Table I: What determines cycle frequency?

	Frequency determined by	**Cross-industry prediction**
Canonical RBC	Shock autocorrelation	Common if shocks are common
Canonical NK	Policy rule	Common across industries
Beaudry et al. (2020)	Bifurcation condition	Not the object of their analysis
Timing search	$\boldsymbol{\omega = \sqrt{\varphi/\mu}}$	**Industry-specific frequency**

Notes: In canonical RBC and NK specifications, the spectral peak reflects the shock process or the policy rule—objects common across industries; multi-sector enrichments can produce dispersion but do not, by themselves, impose a γ-dependent slope (Online Appendix F.2). In Beaudry et al. (2020), the frequency emerges from a Hopf bifurcation; cross-industry frequency variation is not the object of their analysis. In timing search, the frequency is a closed-form function of two market-structure primitives (φ: profit-landscape curvature; μ: timing friction), is exact under linear utility, and is isolated from the shock process (Proposition 3).

[5] This superposition is a microfounded Fourier decomposition of the aggregate cycle. The ironing condition pins each unit's optimal entry path as the harmonic function $n_i(t) = n_{p,i} + A_i \cos(\omega_i t + \phi_i)$ at its own structural natural frequency $\omega_i = \sqrt{\varphi_i/\mu_i}$—the individual optimal decision is itself a Fourier basis element by construction, not a basis element imposed from outside. Pushed to a continuous heterogeneity index, the aggregate cycle $X(t) = \int n_i(t)\, dF(i)$ is the literal superposition of these basis elements weighted by the heterogeneity distribution F over (φ, μ), with bandwidth set by bottleneck damping (Section 4.4). The decomposition relies on three conditions: idiosyncratic shocks, no input-output linkages, and linear utility. With granular establishment- or region-level data, the empirical aggregate spectrum can be read as a structural deconvolution problem on F, an extension that current 3-digit cross-sectional tests do not exploit.

3.5 Intuition: why is the frequency $\sqrt{\varphi/\mu}$?

The closed-form result is the timing-search industry behaving like a spring.

Why the industry oscillates: a spring with stiffness φ and inertia μ. As an industry expands toward optimal scale, firms keep entering because profits are above the steady-state norm. The ironing condition—no firm can profit by entering at a different date—pins entry to a smooth, equalized rate along the approach. That pinned rate has a sharp implication: at n_p, entry cannot drop discontinuously to zero without creating arbitrage, so the industry overshoots. Beyond n_p, the same indifference logic drives entry to the other side of replacement, the flow reverses, and the industry oscillates rather than settling. The economics maps cleanly onto classical mechanics: profit curvature φ acts as a *restoring force* (steeper fall, harder pull back); the timing friction μ acts as *inertia* (larger μ, harder to change the pace of entry). A mass-spring system at $\omega = \sqrt{k/m}$ oscillates at the same frequency for the same reason as the timing-search industry at $\omega = \sqrt{\varphi/\mu}$. Both forces are necessary: take away the scale externality and the spring has no stiffness; take away the matching friction and the industry has no inertia, and oscillation collapses into a quiet return to optimum.

Isochrony: big and small cycles run on the same clock. The cycle's period does not depend on how hard the industry is hit. A deep recession leaves the industry far below optimal scale, profits per active firm are well above norm, and the ironing condition sustains a fast inflow; but the deeper trough means the industry has further to travel. A small shock generates weaker incentive, slower entry, and a shorter path. Speed and distance scale together, so $T = 2\pi\sqrt{\mu/\varphi}$ is the same in either case. Isochrony extends to the damped equilibrium: successive swings preserve their period even as the envelope decays (Online Appendix C.1). A direct corollary is the plucking pattern: deeper recessions recover faster (maximum return velocity $A\omega$ scales linearly with amplitude) but in the same time, the timing-search analogue of the plucking pattern documented by Friedman (1993), Bordo and Haubrich (2017), and Dupraz et al. (2019), and tested in Section 5.4.

3.6 Properties of the cycle

Corollary 1 (Properties of the deterministic cycle)**.** *Under Theorem 1:*

(i) Comparative statics. $\partial\omega/\partial\varphi > 0$ *and* $\partial\omega/\partial\mu < 0$.

(ii) Period. $T = 2\pi\sqrt{\mu/\varphi}$, *depending only on* (φ, μ).

(iii) Amplitude invariance. *ω is independent of the cycle amplitude A and of productivity z, which scales profits without changing curvature.*

(iv) Free amplitude. *The deterministic system is a center: orbits form a one-parameter family indexed by the equilibrium value level E, and the deterministic dynamics do not select among them. The stochastic extension (Theorem 2) selects the unique stationary* equilibrium.

Proof. Items (i)–(iii) follow by direct calculation from $\omega = \sqrt{\varphi/\mu}$ in Theorem 1, which contains no dependence on (z, A). Item (iv) is the energy-conservation property of the frictionless law of motion (13): in the absence of dissipation, $\frac{\mu}{2}\dot{n}^2 + \frac{\varphi}{2}(n - n_p)^2 = E$ is conserved along orbits, so each value of $E \geq 0$ traces a distinct closed orbit and the dynamics fix the period but not the amplitude. ■

Beyond the Gaussian The closed-form result extends to general hump-shaped profits: under any C^2 profit function π with $\pi'(n_p) = 0$ and $\pi''(n_p) < 0$, the frequency takes the form $\omega = \sqrt{|\pi''(n_p)|/[\mu\,\pi(n_p)]} + O(A^2)$, and the Gaussian-based prediction $T_0 = 2\pi\sqrt{\mu/\varphi}$ stays within 1.6% of the exact orbit-integrated period at $A/n_p = 0.30$ across 8 smooth hump-shaped profit specifications (Online Appendix B.3 and C.2). Exact isochrony is the $c_4 = 0$ knife-edge of the finite-amplitude expansion $\omega(A) = \omega_0(1 + c_4 A^2 + O(A^4))$, which the Gaussian satisfies (Online Appendix B.5). Relaxing the Gaussian's exact symmetry generates the empirically standard asymmetry in expansions and contractions: when $\pi'''(n_p) < 0$, the economy spends more time below n_p (slow expansions, sharp contractions), with the asymmetry parameter pinned by $\pi'''(n_p)/|\pi''(n_p)|$ (Online Appendix B.5).

3.7 Balanced growth cycle

The closed-form result is stated around a stationary n_p, but the model accommodates balanced growth at the level of primitives provided three homogeneity conditions hold: intensive-form profit $\pi(n, t) = z(t)\hat{\pi}(n/n_p(t))$ with $z(t)$ and $n_p(t)$ trending at exogenous rate g and $\hat{\pi}$ hump-shaped around 1; matching technology $m(S, R) = S^\alpha R^{1-\alpha}$ homogeneous of degree 1 (already satisfied); and timing friction in percentage rates, $\xi(\dot{n}/n) = \exp(-\frac{\mu}{2}(\dot{n}/n)^2)$.[6] The household discount rate satisfies $\rho > g$.

Proposition 2 (Balanced growth cycle)**.**
Under the homogeneity setup above, the timing search equilibrium of Theorem 1 extends

[6]The body's $\xi(\dot{n})$ coincides with $\xi(\dot{n}/n)$ at leading order around the stationary n_p; the percentage-rate version is the BGP-compatible primitive specification.

to balanced growth: the equilibrium path decomposes as $\ln n(t) = \ln n_p(t) + \tilde{x}(t)$*, where the detrended log deviation* $\tilde{x}(t)$ *obeys*

$$\mu\,\ddot{\tilde{x}} \;+\; \varphi\,\tilde{x} \;=\; 0, \qquad \tilde{x}(t) = \tilde{A}\cos(\omega t + \phi), \qquad \omega = \sqrt{\varphi/\mu}, \tag{16}$$

with $\varphi \equiv -\hat{\pi}''(1)/\hat{\pi}(1)$ *the log-curvature of intensive-form profit at the optimum.*

Proof. Detrending the agent's value function with effective discount $\rho - g > 0$ yields a stationary problem in $\tilde{n}$; the ironing condition becomes $\hat{\pi}(\tilde{n})\,\xi(\dot{\tilde{n}}/\tilde{n}) = \tilde{C}$, and linearization around $\tilde{n} = 1$ delivers (16). ■

Two implications. The frequency $\omega = \sqrt{\varphi/\mu}$ is invariant to the BGP rate g, so the cross-industry test in Section 5.3 extends to industries on different growth trends; and growth-side (g) and cycle-side (φ, μ, χ) policy interventions are structurally separable in the model, complementing the entry-side taxonomy of Section 6.2 with a growth-side dimension.

4 Bottleneck friction and the stationary cycle

The frictionless cycle of Theorem 1 is a center: a one-parameter family of closed orbits with free amplitude. Two ingredients are needed to convert this skeleton into the equilibrium that confronts the data, and each contributes a distinct empirical signature. *Bottleneck friction* (Section 4.1) on its own dissipates cycle amplitude over time while preserving the structural period—the plucking / manufacturing-echo signature documented in Section 5.4. *Stochastic shocks* (Section 4.3) on their own preserve the spectral peak at the structural frequency $\omega_0 = \sqrt{\varphi/\mu}$, but the amplitude diffuses without bound. Together (Section 4.4) they yield the damped stochastic oscillator at stationary amplitude $\sigma_x = \sigma_\varepsilon/\sqrt{2\chi\varphi}$. The Selection Theorem of Section 4.5 then sharpens this construction in two ways: it shows that the resulting stationary distribution is the *unique* Markov stationary timing-search equilibrium, and it isolates the small-noise limit at which the deterministic benchmark re-emerges sharply. Section 4.6 closes with the frequency-welfare separation that the stationary cycle delivers in closed form.

4.1 Bottleneck friction

Assumption 4 (Bottleneck friction)**.**
The supplier-side resource flow contributes an additive quadratic congestion cost $\frac{1}{2}\chi\,\dot{n}^2$ *with*

$\chi > 0$, *generating a damping force* $-\chi\dot{n}$ *in the equilibrium equation of motion via the quadratic dissipation function* $\mathcal{R}(\dot{n}) = \frac{1}{2}\chi\dot{n}^2$.[7]

What the bottleneck friction captures, and how it differs from μ. Although both summarize congestion, μ and χ sit on opposite sides of the entry process and play different roles in equilibrium. The matching friction μ is on the *firm side*—the matching cost each entering firm pays inside the kernel $\xi(\dot{n}) = e^{-(\mu/2)\dot{n}^2}$, microfounded as $\mu = (1-\alpha)/[\alpha(\delta n_p)^2]$ from Cobb-Douglas matching (§2.3)—and acts as inertia in the equilibrium law of motion, pinning the cycle's frequency. The bottleneck friction χ is on the *supplier side*—the convex resource cost the supplier sector pays to deliver net adjustment $\dot{n}$ at peak strain, microfounded as $\chi = c_2 r^2$ from the bottleneck curvature c_2 and the resource-flow linkage r (Online Appendix B.2)—and acts as dissipation, pinning the cycle's amplitude. The disjoint determinants of frequency ($\omega_0 = \sqrt{\varphi/\mu}$) and welfare ($\mathbb{E}[\mathcal{L}] = \sigma_\varepsilon^2/(2\chi)$) in Theorem 3 follow directly from this separation of mechanism: μ is paid inside firms' profits, χ is paid by households for supplier-side infrastructure.

4.2 Deterministic damped case: amplitude decays, period stays

Before adding shocks, bottleneck friction acting alone already delivers a distinctive cycle property that the data can verify. Setting $\sigma_\varepsilon = 0$ in the equilibrium law of motion with bottleneck friction:

$$\mu\,\ddot{n} \;+\; \chi\,\dot{n} \;+\; \varphi(n - n_p) \;=\; 0. \tag{17}$$

For an industry displaced from n_p by a one-time shock—an initial deviation A_0 with zero initial velocity—the equilibrium path is

$$n(t) - n_p \;=\; A_0\, e^{-\eta t/2}\Big[\cos(\omega_d\, t) + \tfrac{\eta}{2\omega_d}\sin(\omega_d\, t)\Big], \qquad \omega_d \equiv \sqrt{\omega_0^2 - \eta^2/4}, \quad \eta \equiv \chi/\mu,$$

where the speed at which bottleneck friction dissipates output is $\eta = \chi/\mu$, the ratio of bottleneck friction to the matching friction's inertial weight. Three properties follow.

- *The period barely changes.* At empirical damping ($\eta^2/(4\omega_0^2) \lesssim 5\%$) the cycle length stays essentially at the structural $T = 2\pi/\omega_0$ throughout the decay, rising at most $O(\eta^2)$ above it (since $\omega_d = \sqrt{\omega_0^2 - \eta^2/4} < \omega_0$).
- *The amplitude shrinks toward steady state.* Each successive swing is smaller than the

[7] This dissipation function is known in classical mechanics as the Rayleigh dissipation function. I adopt "bottleneck friction" for χ throughout the body to signal its economic role as a supplier-side resource bottleneck in the entry pipeline, and reserve "Rayleigh" for the formal derivation in Online Appendix B.2.

last by a factor that depends on the ratio of bottleneck friction χ to matching friction μ; the amplitude does not diverge.

- *Period and amplitude decay are independent.* A heavier bottleneck shrinks amplitude faster but barely changes the cycle length, in contrast to AR(1) and similar linear-decay models where persistence and period are mechanically linked.

This is the *plucking / echo pattern*: a one-time shock triggers oscillations that lose amplitude over time while keeping their period. U.S. manufacturing investment after the Great Depression onset oscillates at $T = 5.0$ years through the early 1970s while the amplitude decays (Section 5.4); oil and mining investment after 1979 oscillates at $T = 3.25$ years. The pattern distinguishes bottleneck dissipation from linear-decay alternatives such as AR(1), where decay rate and period are mechanically linked (Online Appendix C.1 plots the impulse responses).

4.3 Adding shocks

Consider an i.i.d. disturbance to profits:

$$\tilde{\pi}(n,t) = \pi(n) + \sigma_\varepsilon \, \varepsilon_t, \qquad \varepsilon_t \overset{\text{i.i.d.}}{\sim} (0,1). \tag{18}$$

The ironing condition no longer holds exactly: the equilibrium value level E drifts stochastically with the shock, and the economy migrates between orbits. Without bottleneck friction this migration is unbounded—the cycle's amplitude grows without bound on average and no stationary distribution exists—but the spectral peak remains at the structural frequency $\omega_0 = \sqrt{\varphi/\mu}$.

Proposition 3 (Stochastic limit cycle)**.**
Under small shocks (18)*:*

(i) The spectral density of $n(t)$ peaks at the structural frequency $\omega_0 = \sqrt{\varphi/\mu}$ in the undamped small-shock limit; with bottleneck damping $\eta = \chi/\mu > 0$, the observed PSD peak is $\omega_{PSD} = \sqrt{\varphi/\mu - \eta^2/2}$, equal to ω_0 up to an $O(\eta^2)$ correction (Proposition 4(iv)).

(ii) The peak frequency is determined by (φ, μ) at leading order and by the damping ratio η at second order, not by σ_ε.

(iii) Even i.i.d. shocks produce persistent cycles: the persistence is endogenous to the equilibrium structure.

Proof. See Online Appendix B.1. ■

The structural frequency is a property of market structure, not the shock process: i.i.d. shocks produce a PSD peak at $\omega_{\text{PSD}} = \sqrt{\omega_0^2 - \eta^2/2}$, while persistent shocks shift the peak negligibly when the frequency response is sharp. This matches the central finding of Beaudry et al. (2020)—persistent boom-bust dynamics from short-lived disturbances—though the mechanism here is the scale-congestion trade-off rather than financial-friction complementarities; numerical confirmation is in Online Appendix C.1.

4.4 The stationary cycle: damped stochastic oscillator

In the stationary cycle—with both bottleneck dissipation and profit shocks—the industry oscillates at the same natural frequency $\omega = \sqrt{\varphi/\mu}$ but with a well-defined stationary amplitude: dissipation pulls the industry toward optimal scale, shocks re-excite it, and the balance pins the variance of $(n - n_p)$ at $\sigma_\varepsilon^2/(2\chi\varphi)$ (Proposition 4). Differentiating the ironing condition and balancing the value-erosion rate against the bottleneck's output-dissipation rate $\chi\dot{n}^2$ yields the equation of motion (Online Appendix B.2):

$$\underbrace{\ddot{n}}_{\substack{\text{change in}\\ \text{entry rate}}} + \underbrace{\eta\,\dot{n}}_{\substack{\text{congestion}\\ \text{damping}}} + \underbrace{\frac{\varphi}{\mu}(n - n_p)}_{\substack{\text{pull toward}\\ \text{optimal scale}}} = \underbrace{\frac{\sigma_\varepsilon}{\mu}\varepsilon_t}_{\substack{\text{productivity}\\ \text{shock}}}, \tag{19}$$

with $\eta = \chi/\mu$. The natural frequency $\omega_0 = \sqrt{\varphi/\mu}$ is unchanged: bottleneck friction and shocks shape how severely the industry oscillates, not the period at which it does so.

Proposition 4 (Amplitude selection)**.**
With bottleneck friction coefficient $\chi > 0$ (damping $\eta = \chi/\mu > 0$) and shocks $\sigma_\varepsilon > 0$:

(i) Without shocks ($\sigma_\varepsilon = 0$): orbits spiral inward to $(n_p, 0)$ with decay half-life $\ln 2 \cdot 2\mu/\chi$.

(ii) With shocks: the system has a unique stationary distribution with variance:

$$\mathbb{E}\big[(n - n_p)^2\big] = \frac{\sigma_\varepsilon^2}{2\,\eta\,\varphi\,\mu} = \frac{\sigma_\varepsilon^2}{2\,\chi\,\varphi}, \tag{20}$$

where the second equality uses $\eta = \chi/\mu$, so that the matching friction μ cancels.

(iii) The amplitude envelope $A = \sqrt{(n - n_p)^2 + (\dot{n}/\omega_0)^2}$ follows a Rayleigh distribution with scale $\sigma_x = \sigma_\varepsilon/\sqrt{2\chi\varphi}$, so that $\mathbb{E}[A^2] = 2\sigma_x^2 = \sigma_\varepsilon^2/(\chi\varphi)$.

(iv) The spectral peak frequency of the driven oscillator's stationary spectrum is

$$\omega_{peak} = \sqrt{\varphi/\mu - \eta^2/2}.$$

It is obtained by maximizing

$$|H(\omega)|^2 = \frac{1}{(\omega_0^2 - \omega^2)^2 + \eta^2\omega^2}$$

over ω*. The homogeneous damped natural frequency* $\sqrt{\varphi/\mu - \eta^2/4}$*, at which the unforced solution oscillates, is a separate object. For small* η*, both quantities are negligibly below* $\sqrt{\varphi/\mu}$*.*

Proof. See Online Appendix B.2. ■

Figure 1 visualizes the four dissipation/shock regimes: dissipation alone makes orbits spiral inward to n_p (top right), shocks alone make orbits diffuse without bound across amplitudes (bottom left), and only with both ingredients (bottom right) does the balance pin a stationary distribution at $\mathbb{E}[(n - n_p)^2] = \sigma_\varepsilon^2/(2\chi\varphi)$. Online Appendix C.1 confirms numerically that the amplitude envelope follows the Rayleigh density of Proposition 4(iii).

Substituting $\eta = \chi/\mu$ gives $\sigma_x = \sigma_\varepsilon/\sqrt{2\chi\varphi}$, in which μ cancels exactly: industries with volatile profits, low bottleneck friction, or flat profit landscapes experience larger cycles, while μ governs only the cycle frequency. Combining with the welfare formula (24):

$$\mathbb{E}[\mathcal{L}] \approx \frac{\varphi}{2}\,\mathbb{E}[A^2] = \frac{\sigma_\varepsilon^2}{2\,\chi}. \tag{21}$$

Higher bottleneck friction χ *reduces* welfare loss by damping the cycle—a "stabilization through friction" result—with the welfare cost depending only on $(\sigma_\varepsilon, \chi)$ as φ cancels.

4.5 Selection of the stationary equilibrium

Proposition 4 pins the stationary variance of $(n - n_p)$ under the ironing condition. The next theorem strengthens that result in three ways. Under positive bottleneck friction and positive shocks, the stochastic dynamics of (9) admit a *unique* Markov stationary timing-search equilibrium in the sense of Definition 2, and this equilibrium is globally attracting from every initial law. As shocks vanish, the deterministic equilibrium of Theorem 1 re-emerges sharply as the limit.

Theorem 2 (Selection of the stationary equilibrium)**.**
Under Assumptions 1–4 and $\sigma_\varepsilon, \chi > 0$*, there exists a unique Markov stationary timing-search equilibrium* (μ^*, e^*, V^*) *in the sense of Definition 2. The stationary distribution* μ^* *is Gaussian*

Figure 1: Amplitude selection through bottleneck dissipation

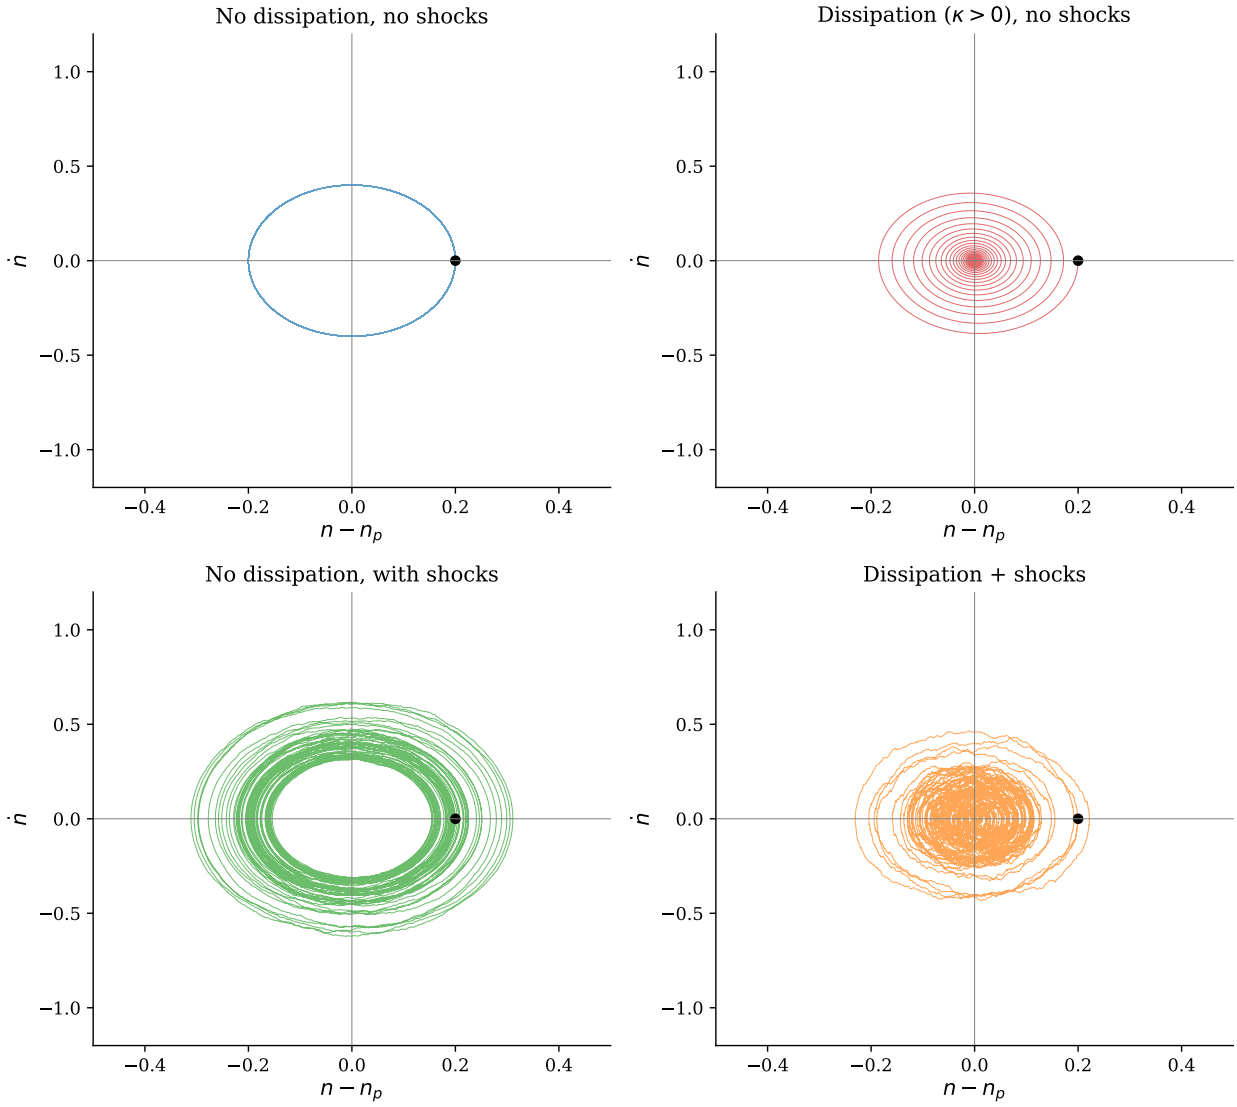

Notes: Phase portraits $(n - n_p, \dot{n})$ under four regimes. Top left: deterministic orbit (closed ellipse, free amplitude). Top right: dissipation without shocks (inward spiral to steady state). Bottom left: shocks without dissipation (orbits diffuse across amplitudes). Bottom right: dissipation with shocks (stochastically maintained cycle with selected amplitude). The balance of dissipation and forcing pins $\mathbb{E}[(n - n_p)^2] = \sigma_\varepsilon^2/(2\eta\varphi\mu)$.

on $\mathbb{R}^2$ *with mean* $(n_p, 0)$ *and covariance*

$$\Sigma^* = \frac{\sigma_\varepsilon^2}{2\,\chi\,\varphi} \begin{pmatrix} 1 & 0 \\ 0 & \varphi/\mu \end{pmatrix},$$

and is globally exponentially attracting: starting from any initial distribution ν_0 *with finite*

second moments, the distribution of $(n_t, \dot{n}_t)$ *converges to* μ^* *at exponential rate,*

$$\left\|\mathcal{L}\big((n_t, \dot{n}_t) \mid \nu_0\big) - \mu^*\right\|_{\mathrm{TV}} \;\le\; C\, e^{-\lambda t}, \qquad t \ge 0,$$

for some $C, \lambda > 0$ *that depend only on the primitives. Selection is sharp at two boundaries. Along* $\sigma_\varepsilon \downarrow 0$ *with* χ *fixed,* μ^* *concentrates on the deterministic steady state of Theorem 1. Along* $\sigma_\varepsilon, \chi \downarrow 0$ *jointly with* $\sigma_\varepsilon^2/\chi \to 2\bar{T}$*,* μ^* *converges to the unique stationary distribution at expected per-firm welfare cost* $\bar{T}$*.*

Proof. See Online Appendix B.3. The proof has four steps: existence with explicit covariance via the matrix equation pinning the stationary covariance; uniqueness from standard linear-Gaussian arguments; exponential mixing in total variation; and the two vanishing-noise selection limits via standard small-noise exit estimates. ■

Theorem 2 supplies three things that Propositions 3–4 did not. First, *uniqueness* of the stationary distribution as an *equilibrium* object, not merely as the long-run limit of the stochastic dynamics; clause (iii) of Definition 2 is the binding restriction. Second, a sharp deterministic-limit correspondence that resolves the free-amplitude indeterminacy of Theorem 1 from outside the conservative landscape. Third, a small-noise selection rule that picks out a unique stationary distribution at any pre-specified expected per-firm welfare cost; Section 4.6 uses this rule to derive Theorem 3 in two lines.

4.6 Frequency-welfare separation

The closed-form structural frequency $\omega_0 = \sqrt{\varphi/\mu}$ (Theorem 1), the stationary variance $\sigma_x^2 = \sigma_\varepsilon^2/(2\chi\varphi)$ (Proposition 4 after substituting $\eta = \chi/\mu$, so that μ cancels), and the welfare cost $\mathbb{E}[\mathcal{L}] = \sigma_\varepsilon^2/(2\chi)$ (Section 6.1) together establish the paper's central structural claim: *structural frequency and welfare cost have disjoint determinants*, with cycle severity standing as an intermediate observable depending on the overlapping set $(\sigma_\varepsilon, \chi, \varphi)$. I state this as a theorem.

Theorem 3 (Frequency-welfare separation)**.**
In the timing search equilibrium under Assumptions 1–4 and the damped-stochastic specification of Proposition 4, with $\eta = \chi/\mu$*:*

$$\begin{aligned}
&\textit{structural frequency:} && \omega_0 = \sqrt{\varphi/\mu} && \textit{depends only on } (\varphi, \mu);\\
&\textit{cycle severity:} && \sigma_x = \sigma_\varepsilon/\sqrt{2\,\chi\,\varphi} && \textit{depends on } (\sigma_\varepsilon, \chi, \varphi);\\
&\textit{welfare cost:} && \mathbb{E}[\mathcal{L}] = \sigma_\varepsilon^2/(2\chi) && \textit{depends only on } (\sigma_\varepsilon, \chi).
\end{aligned}$$

The structural frequency depends on (φ, μ); the welfare cost depends on $(\sigma_\varepsilon, \chi)$; these determinant sets are disjoint, with cycle severity the intermediate observable connecting both. The structural frequency is invariant to σ_ε, χ, the shock persistence (in the i.i.d. specification), and the realized cycle amplitude; the welfare cost is invariant to (φ, μ) and to the per-firm entry cost κ. The observed PSD peak inherits the structural frequency up to an $O(\eta^2)$ correction: $\omega_{PSD} = \sqrt{\varphi/\mu - \eta^2/2}$ (Proposition 4(iv)).

Proof. $\omega_0 = \sqrt{\varphi/\mu}$ is the natural frequency of the deterministic value landscape (Theorem 1); it depends only on (φ, μ). $\mathbb{E}[\mathcal{L}^*] = \sigma_\varepsilon^2/(2\chi)$ is the expected per-firm cycle-induced welfare cost under the unique Markov stationary equilibrium μ^* (Theorem 2 combined with the welfare formula (21)); it depends only on $(\sigma_\varepsilon, \chi)$. The two are therefore pinned by disjoint coordinates of the primitive vector $(\mu, \chi, \varphi, \sigma_\varepsilon)$ by construction. The cycle-severity formula $\sigma_x = \sigma_\varepsilon/\sqrt{2\chi\varphi}$ is Proposition 4(ii) under the substitution $\eta = \chi/\mu$, and the $O(\eta^2)$ PSD-peak correction is Proposition 4(iv). Disjointness is preserved at the small-noise boundary: under the joint scaling of Theorem 2 ($\sigma_\varepsilon, \chi \downarrow 0$ with $\sigma_\varepsilon^2/\chi \to 2\bar{T}$), the welfare cost limits to $\bar{T}$ while the natural frequency is unchanged. ■

The theorem has three methodological consequences. First, cross-industry variation in ω_0 identifies the market-structure ratio φ/μ free of contamination from χ, σ_ε, or κ (the $O(\eta^2)$ PSD-peak shift is below 5% at empirical damping); the cross-sectional regression of Section 5.3 therefore tests a structural prediction about φ/μ even when proxies for the individual primitives are imperfect. Second, two industries with identical (φ, μ) but different shock environments cycle at the same frequency, making the cross-industry frequency prediction structural rather than a reduced-form artifact of common shocks. Third, industrial policy targeting μ (entry regulation) or φ (cluster development) affects timing alone; only policies that move σ_ε or χ affect welfare (Section 6.2).

Scope of the separation Theorem 3's welfare claim is exact under linear utility, the per-firm effective-output welfare metric of (24), and exclusion of the supplier's $\frac{1}{2}\chi\dot{n}^2$ resource flow; relaxing each condition reintroduces leakage of (φ, μ) into the welfare cost of order 10–25% at the median under standard CRRA preferences or alternative metrics, documented in Online Appendix E.1.

Comparison to alternative endogenous-cycle models Timing search and Beaudry et al. (2020) share the broader insight that endogenous structure can turn short-lived disturbances into persistent cyclical adjustment, but Beaudry et al. (2020) generate a stochastic limit cycle through financial frictions and a Hopf bifurcation, while timing search generates characteristic

frequencies from the entry-date indifference condition and a scale–congestion trade-off, which delivers the industry-decomposable closed form $\omega_0 = \sqrt{\varphi/\mu}$ tested in Section 5.3. Across standard benchmarks (AR(1) RBC, NK AR(2), multi-sector RBC with input-output linkages, and Khan and Thomas (2008)-style heterogeneous-θ adjustment costs) the simulated cross-industry slope $\hat{\beta}_\gamma$ centers near zero with 95% CIs roughly $[-0.13, +0.20]$, none of 50 replications reaching the empirical $[0.50, 0.75]$ range; the comparison clarifies what additional structure the empirical question requires rather than ruling out enriched versions of these models (Online Appendices F.1–F.4).

5 Model vs. data

U.S. industries fluctuate at substantially different frequencies. Across six broad industry aggregates with bell-shaped interior peaks within the business-cycle band of 2–12 years (Burg AR(24) on monthly industrial production, 1972–2024), spectral peak periods range from 3.0 years (Nondurable Manufacturing) to 4.7 years (Business Equipment) (Section 5.2, Figure 3)—a more than 50% spread across industries that share the same monetary policy, aggregate demand shocks, and financial conditions. Models that derive spectral properties from a common shock process or policy rule have difficulty reconciling this dispersion. Theorem 1's closed form $\omega = \sqrt{\varphi/\mu}$ predicts that the variation is *structural*: each industry's cycle frequency is pinned by market-structure primitives, not by the shock process. The remainder of this section organizes the empirical evidence around five questions the timing-search mechanism implies: whether industries cycle at different frequencies (Section 5.2); whether the structural slope $\beta_\gamma = 1/2$ holds (Section 5.3, the headline test); and three further diagnostics on whether the dynamics are genuinely oscillatory, whether the structural separation theorem holds across industries, and whether the model's predictions about cycle shape and recovery hold (Section 5.4).

5.1 Cross-industry predictions

Taking logs of $\omega_i = \sqrt{\varphi_i/\mu_i}$ yields the cross-industry regression framework

$$\ln \omega_i = \tfrac{1}{2} \ln \varphi_i - \tfrac{1}{2} \ln \mu_i + u_i, \tag{22}$$

with the slopes $\pm 1/2$ as exact model predictions. The headline test of Section 5.3 restricts attention to 15 manufacturing 3-digit industries, within which the matching technology (permits, specialist construction labor, equipment supply chains) is approximately common across industries and the turnover rate $\delta_i \in [0.065, 0.090]$ varies only by a factor of 1.4. Under

approximate μ-uniformity, the prediction reduces to a φ-driven slope $\beta_\gamma = 1/2$. Because φ_i is unobserved, the empirical implementation substitutes the geographic-concentration index γ_i under a maintained measurement equation $\ln \varphi_i = a + \lambda \ln \gamma_i + v_i$, so the regression slope identifies $\lambda/2$, equal to 1/2 under unit log-elasticity ($\lambda \approx 1$); the evidence supporting this restriction is collected in Section 5.3. Two industries facing identical shock processes but differing in (φ, μ) are predicted to cycle at distinct frequencies; this prediction distinguishes the timing-search framework from RBC and New Keynesian frameworks, in which the spectral peak inherits its location from the shock process or the policy rule.

5.2 Spectral evidence from U.S. industry data

Do industries cycle at different frequencies? I evaluate the cross-industry prediction using monthly industrial production indices from FRED for eight U.S. sectors over 1972–2024. The series span a range of market structures: from Motor Vehicles & Parts (dense supplier networks, strong scale externalities) to Mining (dispersed extraction, weak inter-firm linkages) to Electric Power Generation (regulated, capital-intensive).

Computing spectral peaks via Burg AR(24) on monthly log-differenced production for each industry over 1972–2024 reveals substantial cross-industry variation within the business-cycle band of 2–12 years. Six of the eight FRED aggregates have interior bell-shaped peaks within the band (Figure 2): Business Equipment has the longest peak period at 4.7 years; Manufacturing aggregates and Mining cluster near 3.5–4.0 years (Total Industrial Production 3.9, Manufacturing 4.0, Durable Manufacturing 3.9, Mining 3.6); Nondurable Manufacturing peaks at 3.0 years. Two industries (Motor Vehicles & Parts and Electric Power Generation) concentrate spectral mass at the high-frequency boundary of the band, indicating sub-business-cycle dynamics, and are excluded from the cross-industry summary below. The dispersion across the six interior-peak industries is consistent with ω depending on market structure rather than on common shocks.

Figure 3 summarizes the dominant cycle period for each of the six interior-peak industries. The cross-industry dispersion—ranging from 3.0 years to 4.7 years across industries that share a common macroeconomic environment—is difficult to reconcile with models that derive spectral properties from a common shock process. In the timing search framework, the variation is structural: it reflects cross-industry differences in market-structure primitives. Across this six-industry sample—which spans aggregate Manufacturing and Mining—both the profit-landscape curvature (φ_i) and the timing friction (μ_i) plausibly vary across sectors with different matching technologies. The headline test of Section 5.3 narrows to 15 manufacturing 3-digit industries, within which μ is approximately uniform, so the cross-industry signal

Figure 2: Spectral density of U.S. industry production growth (six interior-peak industries)

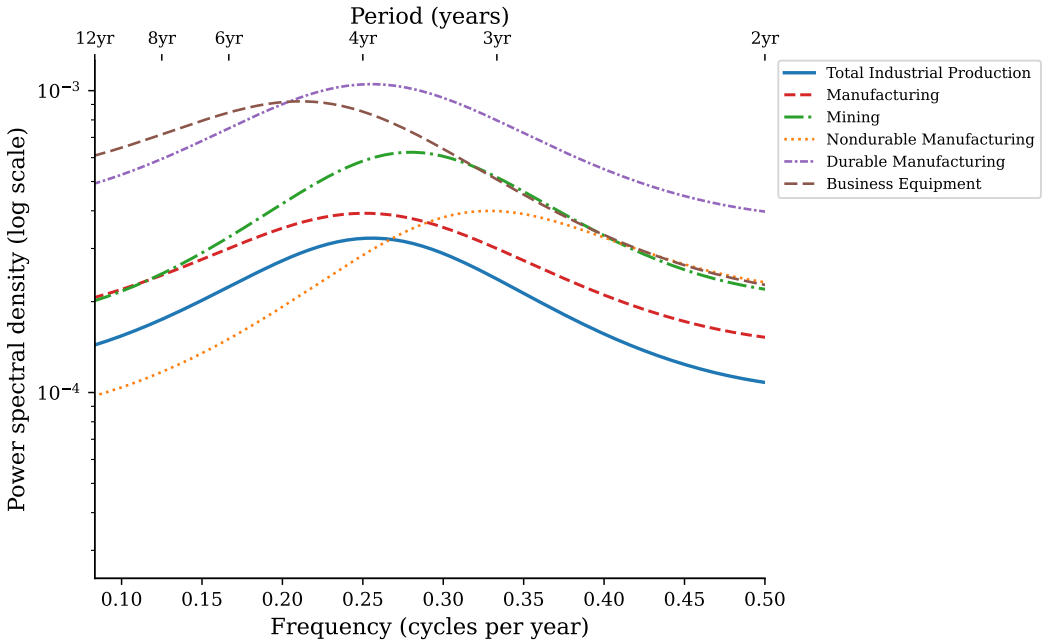

Notes: Burg AR(24) power spectral density of monthly log-differenced industrial production indices from FRED, 1972:M1–2024:M12, restricted to the business-cycle band of 2–12 years. The top axis labels convert frequency (cycles per year) to period (years). Each industry exhibits a bell-shaped interior peak within the band; peak periods range from 3.0 years (Nondurable Manufacturing) to 4.7 years (Business Equipment). Motor Vehicles & Parts and Electric Power Generation are excluded; their spectral mass concentrates at the high-frequency boundary, indicating sub-business-cycle dynamics.

collapses to the φ channel.

Model shape versus data shape The timing-search SDE generates an internal spectral peak at $\omega_0 = \sqrt{\varphi/\mu}$, while an AR(1) on growth rates produces a smooth low-pass spectrum with no internal peak. Figure 4 compares the model PSD at a manufacturing-like calibration ($T_0 = 5.0$ years) to the FRED industrial production PSD for NAICS 337 (Furniture), which peaks at $T = 4.9$ years: the model can produce a shape comparable to the data at a single industry, providing qualitative motivation for the cross-industry slope test below.

Consistency check: nonlinear dependence in industry dynamics As a consistency check rather than a discriminating test, I apply the BDS test (Brock et al., 1996) to AR(1)-filtered industry residuals. The i.i.d. null is rejected at the 1% level for all eight industries, with z-statistics ranging from 16.2 (Durable Manufacturing) to 21.8 (Motor Vehicles & Parts). This indicates nonlinear or nonstationary structure beyond AR(1) in industry dynamics. The pattern is consistent with the timing-search mechanism but also with GARCH, regime switching, and structural breaks; the 1972–2024 sample spans the Great Moderation and

Figure 3: Cross-industry variation in spectral peak periods

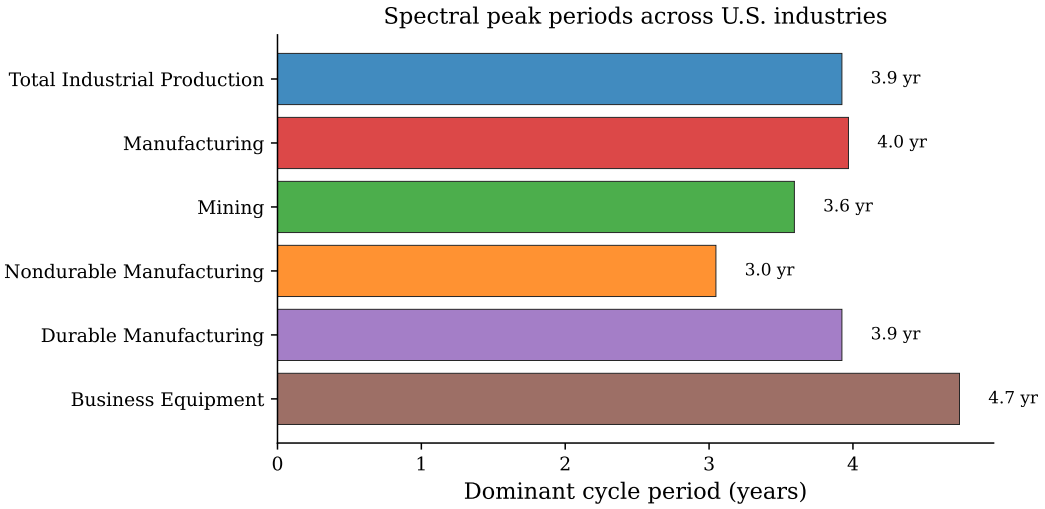

Notes: Dominant cycle period (years) for the six U.S. industries with interior bell-shaped spectral peaks within the business-cycle band of 2–12 years (Burg AR(24) on log-differenced monthly industrial production). Cross-industry peak periods range from 3.0 years (Nondurable Manufacturing) to 4.7 years (Business Equipment); manufacturing aggregates cluster near 3.5–4.0 years. Motor Vehicles & Parts and Electric Power Generation are excluded because their spectral mass concentrates at the high-frequency boundary, indicating sub-business-cycle dynamics.

COVID, both of which can produce comparable BDS rejection magnitudes. The BDS exercise therefore confirms departure from linear AR(1) dynamics but does not discriminate among the nonlinear alternatives. Full results across embedding dimensions 2–6, with AR(1) and AR(p)+GARCH pre-filtering specifications, are in Online Appendix D.2.

5.3 Cross-industry regression with external proxies

Does the structural slope $\beta_\gamma = 1/2$ hold under the maintained restrictions? The cross-industry regression tests the closed-form prediction (22) under three explicit maintained restrictions, each disclosed up front rather than as a footnote: (i) approximate μ-uniformity across 3-digit manufacturing industries, with the observable δ-driven component partialled out as a robustness specification below and Online Appendix D.5; (ii) a measurement equation $\ln \varphi_i = a + \lambda \ln \gamma_i + v_i$ that maps the unobserved profit-landscape curvature to the Ellison and Glaeser (1997) concentration index, so the empirical slope identifies $\lambda/2$, equal to 1/2 at unit log-elasticity ($\lambda = 1$); (iii) the timing-search mechanism's harmonic-oscillator structure dominates at the spectral resolution targeted by the multitaper and Burg AR(12) estimators, which identify the dominant low-frequency mode. Each restriction is either tested directly (i, via the μ-channel three-part defense), disciplined by the revealed-curvature λ diagnostic below while still maintained as an auxiliary link (ii), or transparently dependent on a methodological

Figure 4: Spectral density: timing-search model vs. industry data

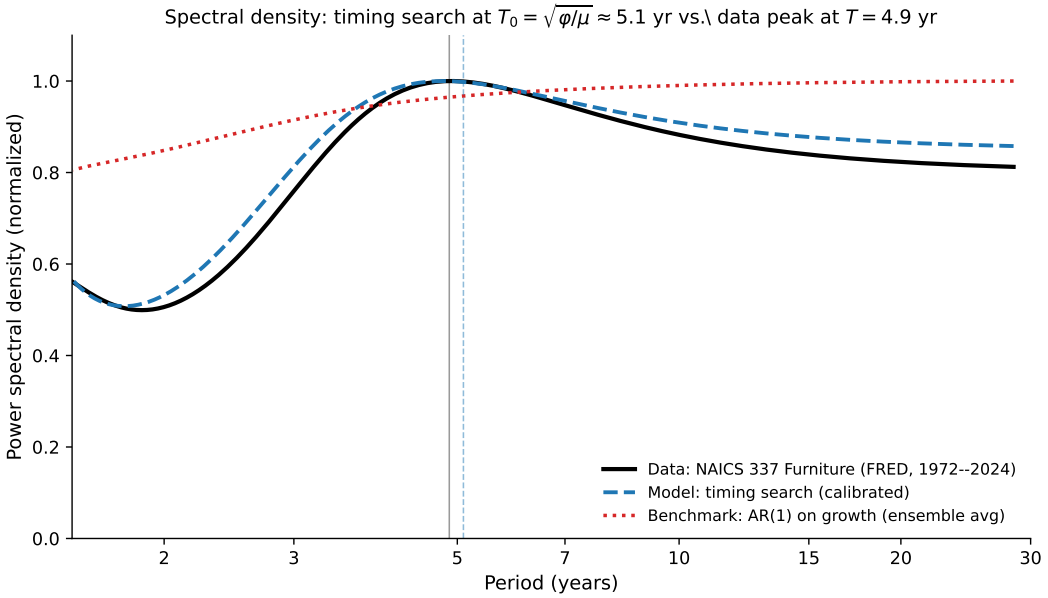

Notes: Burg AR(24) power spectral densities, normalized to the maximum within the business-cycle band (1.5–30 years). Solid black: FRED monthly industrial production growth for NAICS 337 (Furniture), 1972–2024, with empirical peak at $T = 4.9$ years. Dashed blue: simulated timing-search stochastic differential equation at a manufacturing-like calibration ($\varphi = 1.52$, $\mu = 1$, $\chi = 0.32$, $\sigma_\varepsilon = 0.04$, theoretical $T_0 = 2\pi/\sqrt{\varphi/\mu} = 5.1$ years; empirical Burg-AR peak at $T = 4.8$ years), ensemble-averaged over 30 replications, with realistic-conditions noise components (idiosyncratic measurement noise $\sigma_{\text{meas}} = 0.060$ and a slow AR(1) productivity-drift component, $\rho = 0.93$, $\sigma = 0.0011$) that mimic the finite-sample data generating process. Dotted red: ensemble average of 30 AR(1)-on-growth-rate replications ($\rho = 0.5$), shown for comparison—a standard "no internal peak" benchmark with smooth low-pass spectrum. Timing search can generate a comparable peak and surrounding spectral shape; AR(1) cannot.

choice (iii). Under all three, the test is the cross-industry structural-prediction test; outside any of the three, the same regression is a consistency diagnostic for an enriched model with a γ-correlated μ, a non-unit λ, or a broader spectral target.

A baseline regression on seven broad industry groups using BEA proxies (inverse I-O self-purchase ratio for φ and capital per employee for μ) gives a univariate slope of 0.85 ($R^2 = 0.65$, $p = 0.028$) on $\ln \hat{\varphi}_i$ and a multivariate R^2 of 0.75, motivating the broad-sector relationship (Online Appendix D.1). The headline test moves to finer industry resolution: $N = 15$ three-digit manufacturing industries with the Ellison and Glaeser (1997) geographic concentration index γ_i as the φ proxy and multitaper PSD (Thomson DPSS, $K = 5$ tapers) spectral peaks as ω_i, with Burg AR(12) and the 2SLS instrument set as robustness specifications.[8]

Table II reports the slope under several specifications. The cross-industry ω_i measure

[8]The proxy swap is empirical: at the 3-digit level, the broad-sector I-O and capital-intensity proxies do not preserve the relationship; the EG index, constructed at the 3-digit manufacturing level by design, does. Online Appendix D.1 documents.

is the multitaper PSD spectral peak (Thomson DPSS, $K = 5$ tapers) for the headline specifications; column 2 reports a Burg AR(12) robustness check on the spectral side. Across these specifications the empirical slope spans $[0.50, 0.75]$, all within roughly one HC1 standard error of the structural prediction $\beta_\gamma = 1/2$ from (22). Multitaper PSD gives $\hat{\beta}_\gamma = 0.50$ ($p = 0.004$, $R^2 = 0.56$); Burg AR(12) on the same series gives $\hat{\beta}_\gamma = 0.53$ ($p = 0.029$); the strong-first-stage broad-1947 IV gives $\hat{\beta}_\gamma^{\text{IV}} = 0.53$; and the primitive Kim+CBP instrument set gives $\hat{\beta}_\gamma^{\text{IV}} = 0.75$. Wald and Anderson-Rubin tests fail to reject $\beta_\gamma = 1/2$ at conventional levels. The sign matches theory: more geographically concentrated industries—those with stronger agglomeration externalities and higher φ—cycle faster.

Higher-order Burg AR(24). The higher-order Burg AR(24) estimate is lower ($\hat{\beta}_\gamma = 0.24$, $p = 0.004$). The divergence from multitaper and AR(12) is not a finite-sample artifact: AR(1)-, AR(2)-, and BIC-prewhitened Burg AR(24) all give $\hat{\beta}_\gamma \in [0.22, 0.24]$ (Online Appendix D.6). AR(24) over-resolves the BC band, placing the dominant peak at higher frequency than multitaper/AR(12) for long-cycle industries (NAICS 332, 333, 335, 337 shift from $T \approx 6.3$ to ≈ 3.7–4.9 yr), compressing cross-industry ω variance. The structural test is best read off the multitaper and AR(12) specifications, which identify the dominant low-frequency mode the harmonic-oscillator structure maps to.

Historical-geography 2SLS. The main IV is a single 1947 broad-regional specialization score,

$$Z^B_{47,i} = \frac{1}{\sqrt{3}}\Big[z(\ln \text{Hoover}_{47,i}) + z(\ln \text{CBP state HHI}_{47,i}) - z\big(\ln \text{CBP county HHI}_{47,i}\big)\Big],$$

where $z(\cdot)$ standardizes across the 15-industry sample. The score combines Kim's state-level Hoover localization coefficient (Kim, 1995) with state and county concentration from the digitized County Business Patterns of Eckert et al. (2022). The sign convention is fixed by the identifying object: broad regional specialization loads positively, while one-county urban dominance is netted out. Instrumenting modern γ_i with $Z^B_{47,i}$ yields $\hat{\beta}_\gamma^{\text{IV}} = 0.53$ (HC1 SE 0.22, $N = 15$), with first-stage $F_{\text{HC1}} = 17.81$, conventional first-stage $F = 11.32$, and first-stage $R^2 = 0.47$. Standard Wald (1943) inference rejects $\beta_\gamma = 0$ at $p = 0.034$ and fails to reject $\beta_\gamma = 0.5$ at $p = 0.89$; Anderson and Rubin (1949) weak-IV-robust testing gives $p = 0.036$ against zero and $p = 0.87$ against 0.5, and the inverted 95% AR confidence set for β_γ is $[+0.06, +1.09]$ (90% AR: $[+0.18, +0.93]$). The set is bounded and connected—a meaningful first-stage signature, not the unbounded or disjoint set a genuinely weak instrument would produce—but the 95% exclusion of zero is marginal (lower bound 0.06). The honest reading of the AR set is that the IV result is corroborative on sign and order of magnitude rather than a precise test of the structural value: the structural argument runs through the joint

reading of the closed-form prediction (22), the OLS point estimates near 0.5, and the IV's positive bounded set, not through a pointwise IV t-ratio that the $N = 15$ design cannot deliver.

Primitive IV constructions support the same reading without changing the headline setup. Kim 1947 alone gives $\hat{\beta}_\gamma^{\text{IV}} = 0.66$ with first-stage $F_{\text{HC1}} = 5.30$. Adding the 1947 CBP county Herfindahl as a second excluded instrument gives $\hat{\beta}_\gamma^{\text{IV}} = 0.75$ (HC1 SE 0.24), joint first-stage $F_{\text{HC1}} = 7.29$, and Hansen J $p = 0.40$. Adding firm-level 1947 CR4 from the Census of Manufactures (a structurally distinct concentration concept) gives Kim+CR4 joint $\hat{\beta}_\gamma^{\text{IV}} = 0.61$ (HC1 SE 0.16, Hansen J $p = 0.44$); across five IV specifications spanning geographic and firm-level concentration the slope spans $[0.53, 0.75]$ with Hansen J never rejected (Online Appendix D.5). Online Appendix D.5 also reports an instrument-design scan: the broad-1947 score clears $F_{\text{HC1}} > 10$ in the full sample, the strict 2–12-year peak-band sample, and every leave-one-out sample. Some 1927+CBP specifications clear a robust first-stage statistic of 10 only in the full sample but lose strength under the same checks, so I do not use them as the headline first stage.

I use 1947 because, among Kim's three historical years (1900, 1927, 1947), it is the earliest at which all NAICS-3 manufacturing industries are established categories with nonzero Hoover loading; the 1900 coefficient has near-zero first-stage loading on the modern cross-section. The joint 2SLS using all three Kim years yields a similar $\hat{\beta}_\gamma^{\text{IV}} = 0.68$ with overidentification $p = 0.20$.

Partialling out the observable component of μ. The microfounded matching friction $\tilde{\mu}_i = (1-\alpha)/(\alpha\delta_i^2)$ has a measurable component pinned by the turnover rate δ_i, observable from BDS. Under approximate α-uniformity within manufacturing, the structural equation (22) becomes $\ln \omega_i = \text{const} + \frac{1}{2}\lambda \ln \gamma_i + \ln \delta_i + u_i$ (taking the δ-driven part of μ to the left-hand side identifies $\beta_\delta = +1$). I therefore replace "assume μ uniform" with "control for $\ln \delta_i$." Adding $\ln \delta_i$ as an exogenous regressor leaves $\hat{\beta}_\gamma$ essentially unchanged: OLS moves from 0.497 to 0.495 (HC1 SE 0.140 $\to$ 0.149); the broad-1947 IV moves from 0.531 to 0.532 (HC1 SE 0.228 $\to$ 0.231). The estimated $\hat{\beta}_\delta$ is uninformative at $N = 15$ ($\hat{\beta}_\delta = -0.20$ OLS, -0.15 IV; HC1 SE 1.2–1.3): its confidence interval contains both zero and the model's unit-elasticity prediction $\beta_\delta = +1$, reflecting the narrow 1.4× within-manufacturing range of δ_i. The cross-industry slope on $\ln \gamma_i$ is identified off the φ-proxy channel under maintained μ-uniformity; the δ-level prediction is untestable at this resolution rather than contradicted. The headline reading is that the slope on $\ln \gamma_i$ is robust to including the observable μ-component as a control, while the δ-channel itself remains separately unidentified (Online Appendix D.5, "δ-control" specification).

Revealed-curvature diagnostic for λ. The unit-elasticity link $\lambda = 1$ is the weakest auxiliary

step in the quantitative interpretation, so I report a direct diagnostic rather than leaving it implicit. Using the observable turnover component of the matching friction, $\tilde{\mu}_i^{\delta} \propto \delta_i^{-2}$, define a revealed curvature

$$\hat{\varphi}_i^{\delta} \;\propto\; \omega_i^2 \tilde{\mu}_i^{\delta} \;\propto\; \left(\frac{\omega_i}{\delta_i}\right)^2 .$$

This object is not an independent measurement of φ_i because it uses the model's frequency equation and the observed ω_i; it is instead a transparent estimate of the auxiliary measurement elasticity conditional on the same structural equation. Regressing $\ln \hat{\varphi}_i^{\delta}$ on $\ln \gamma_i$ gives $\hat{\lambda} = 1.01$ (HC1 SE 0.28) for the multitaper headline and $\hat{\lambda} = 1.08$ (HC1 SE 0.45) for Burg AR(12). The IV-compatible versions give $\hat{\lambda}^{\mathrm{IV}} = 1.05$ (HC1 SE 0.45) with the broad-1947 instrument and 1.52 (HC1 SE 0.48) with Kim+CBP; none rejects $\lambda = 1$. The same exercise also clarifies what the diagnostic cannot do: an external BEA I-O inverse self-purchase proxy gives a noisy wrong-signed slope ($\hat{\lambda} = -0.22$, HC1 SE 0.23) at NAICS-3 resolution, so it cannot independently identify λ in this sample. The conclusion is therefore a joint-consistency statement: the headline spectral estimators imply a near-unit γ–φ elasticity once the observable δ component of μ is accounted for, but a fully independent estimate of λ would require plant-level productivity-density gradients.

Pre-1947 industrial geography pre-dates the modern macroeconomic environment in which cycle frequencies are observed, motivating the exclusion restriction; I acknowledge the standard caveat that historical concentration could correlate with persistent industry traits that affect modern cycle frequency through channels other than current γ. The most economically substantive form of this caveat—a back-door from 1947 geography through persistent regulation into the matching friction μ—is tested directly in Online Appendix D.5 and the back-door is empirically shut. The IV result is best interpreted alongside the multitaper, AR(12) OLS, and BDS subsector estimates as a coherent body of evidence pointing to β_γ in the neighborhood of the structural prediction 1/2 from (22).[9]

Structural reading of the slope. The model predicts $\omega_i = \sqrt{\varphi_i/\mu_i}$, so the regression slope on $\ln \gamma_i$ in (22) equals 1/2 only if $\ln \mu_i$ is approximately uncorrelated with $\ln \gamma_i$ across the cross-section and the proxy elasticity $\lambda = \partial \ln \varphi / \partial \ln \gamma$ is approximately unity. Within manufacturing, this is plausible on first principles: the matching technology (permits, specialist construction labor, equipment supply chains) is a generic property of the entry-matching pipeline rather than an industry-specific friction, and the structural formula $\mu_i = (1-\alpha)/[\alpha(\delta_i n_{p,i})^2]$, read in percentage-rate units (rescaling $\tilde{n}_i = n_i/n_{p,i}$ removes the scale factor $n_{p,i}^2$), reduces to $\tilde{\mu}_i = (1-\alpha)/(\alpha\delta_i^2)$, which has narrow variation when α is approximately

[9] Robustness to alternative IV constructions and the full three-part defense of μ-uniformity-on-γ are collected in Online Appendix D.5.

uniform and $\delta_i \in [0.065, 0.090]$ (1.4× range across the headline 15 NAICS-3 industries). The empirical recovery of $\beta_\gamma \approx 1/2$ is therefore consistent with μ approximately uniform in scale-normalized units across manufacturing industries: cross-industry frequency variation operates through the profit-landscape curvature φ, which has an order-of-magnitude wider range (γ varies from 0.004 to 0.12, a 30× range).

A direct test of μ-uniformity-on-γ across three external proxies—two capital-side, one regulatory-side—supports the maintained restriction on a three-part defense (Online Appendix D.5). *(i) Capital-side, clean null.* The BEA capital-stock-to-investment replacement horizon and the Cooper and Haltiwanger (2006) quadratic adjustment-cost coefficient ν are statistically indistinguishable from zero on $\ln \gamma_i$ ($\hat{\beta} = +0.02$, $p = 0.85$; $\hat{\beta} = +0.04$, $p = 0.78$). *(ii) Regulatory-side, bounded toward zero.* The QuantGov RegData regulatory restriction count is positive and marginally significant on γ_i ($\hat{\beta} = +0.91$, $p = 0.07$). Under the natural sign convention that the paper's own microfoundation supports—regulation tightens the entry-matching pipeline ("rapid entry overloads permitting and construction"), so regulation enters μ positively—the omitted term $-\frac{1}{2}\ln \mu_i$ in equation (22) biases the OLS slope *toward zero*: the plim is $\beta_\gamma^{\text{OLS}} = \frac{\lambda}{2} - \frac{1}{2}\,\text{cov}(\ln \mu, \ln \gamma)/\text{var}(\ln \gamma)$, with the second term positive. The recovered slope range $[0.50, 0.75]$ is therefore a *lower bound* on the φ-channel slope $\lambda/2$: the structural prediction is met or exceeded despite a contaminant whose sign works *against* finding it. Frisch-Waugh orthogonalization—regressing $\ln \gamma_i$ on $\ln R_i$ and recovering the residual $\tilde{\gamma}_i$, then re-running the regression on $\tilde{\gamma}_i$—confirms slope stability: $\hat{\beta}_{\tilde{\gamma}}^{\text{OLS}} = +0.43$ and $\hat{\beta}_{\tilde{\gamma}}^{\text{IV}} = +0.68$, both within one HC1 SE of 1/2 (Online Appendix D.5, Frisch-Waugh table). *(iii) IV exclusion, defended at the source.* The historical instrument is empirically orthogonal to the modern regulatory contaminant. Regressing $\ln R_i$ directly on the headline composite $Z_{47,i}^B$ yields a coefficient indistinguishable from zero ($\hat{\beta} = +0.20$, HC1 SE 0.91, $p = 0.83$, $R^2 = 0.007$); the same regression on Kim's pure 1947 Hoover component gives $\hat{\beta} = -0.20$ (HC1 SE 1.21, $p = 0.87$, $R^2 = 0.003$). The back-door path 1947 geography → persistent regulation → modern μ → modern ω is empirically shut at the source for both the composite headline IV and its dominant component. The Hansen-J overidentification tests across Kim's Hoover, CBP 1947 Herfindahl, and 1947 CR4 (Online Appendix D.5, Robustness 2–3; J p-values 0.40 and 0.44) further read as a μ-channel exclusion test: if any of the three historical instruments transmitted a regulatory back-door, the J statistic would reject.

Two robustness exercises further support the cross-sectional relationship without overturning it. First, a parallel $N = 10$ cross-section across BDS subsectors with clear interior spectral peaks, using BDS-derived spatial concentration G_i rather than the published γ_i, yields $\hat{\beta} = 0.107$ (HC1 SE 0.043, $p = 0.037$, $R^2 = 0.39$): same sign, similar R^2, different data source for both ω_i and the concentration measure (Online Appendix D.7; sample selection on

Table II: Cross-industry frequency regression

	OLS MT $N = 15$	OLS AR(12) $N = 15$	2SLS Broad $N = 15$	2SLS Kim+CBP $N = 15$	Theory
$\hat{\beta}_\gamma$	0.50	0.53	0.53	0.75	0.50
HC1 SE	(0.14)	(0.22)	(0.22)	(0.24)	—
Wald p vs. 0	0.004	0.029	0.034	0.008	—
Wald p vs. 0.5	0.98	0.89	0.89	0.32	—
AR p vs. 0	—	—	0.036	0.015	—
AR p vs. 0.5	—	—	0.87	0.36	—
Hansen J p	—	—	—	0.40	—
R^2 / 1st-stage F	0.56	0.30	$F = 17.81$	$F = 7.29$	—

Notes: Cross-industry regression of $\ln \omega_i$ on $\ln \gamma_i$ (Ellison and Glaeser 1997 geographic concentration index). Column 1: OLS with multitaper spectral peaks (Thomson DPSS, $K = 5$, Thomson 1982). Column 2: OLS with Burg AR(12) spectral peaks (different ω_i measure as a spectral-side robustness check). Column 3: 2SLS of multitaper ω_i on γ_i using the broad 1947 regional-specialization score Z^B_{47} as a single excluded instrument. Column 4 uses the primitive Kim 1947 Hoover and 1947 CBP county-level employment Herfindahl as two excluded instruments; the reported first-stage statistic is the HC1-Wald joint statistic for the excluded instruments. Spectral peaks estimated from log-differenced monthly industrial production (FRED, 1972–2024). HC1 heteroskedasticity-robust standard errors (White, 1980). AR rows report Anderson and Rubin (1949) weak-IV-robust p-values for the IV columns. The theory column is conditional on the auxiliary unit-elasticity link $\lambda = 1$ in $\ln \varphi_i = a + \lambda \ln \gamma_i + v_i$; the revealed-curvature diagnostic in the text estimates $\hat{\lambda} \approx 1.0$ for the headline multitaper and AR(12) specifications. The broad-1947 score also clears the conventional homoskedastic first-stage threshold ($F = 11.32$) and the strict-band/leave-one-out first-stage checks reported in Online Appendix D.5. The higher-order Burg AR(24) yields $\hat{\beta}_\gamma = 0.24$ (Wald $p = 0.002$ vs. 0.5), the only specification in the Online Appendix D.6 six-estimator panel that rejects 0.5 at the 5% level; the Welch periodogram falls at 1.9 HC1 SE below 0.5 ($p \approx 0.07$) and is reported with AR(24) as a smaller-slope, same-sign sibling.

a feature of the dependent variable). Second, at four-digit resolution ($N = 86$), Ellison et al. (2010) publish γ only at the three-digit level, so every four-digit child inherits its parent's γ and the slope mechanically collapses to $\hat{\beta}_\gamma = 0.027$ ($p = 0.705$); this tests the inheritance of γ rather than its predictive power at four-digit resolution (Online Appendix D.4). An international pooled US+UK exercise (UK ONS, 1990–2024) appears in Online Appendix D.7 with the same-sign result but a documented proxy concern (US γ standing in for UK industry-specific concentration); I do not lean on it in the headline given that BRES-derived UK γ would be required for a clean reading.

Identification The ratio φ/μ is identified directly from the spectral peak ($\varphi/\mu \approx (2\pi/T^*)^2$ at leading order, with an $O(\eta^2)$ correction). Individual φ and μ require external information; χ is identified from $\sigma_x = \sigma_\varepsilon/\sqrt{2\chi\varphi}$ given $(\sigma_\varepsilon, \varphi)$; and κ enters only the steady-state ironing level.

5.4 Further diagnostics

The cross-sectional regression in Section 5.3 is the paper's structural test, identifying the φ channel through the geographic-concentration proxy. I complement it with three further empirical questions, each addressing a distinct prediction of the timing-search mechanism beyond the cross-industry slope. I present these as suggestive rather than definitive evidence; their joint pattern is harder to reproduce in any single alternative model than in any one diagnostic taken alone.

Are the dynamics genuinely oscillatory rather than persistent AR(1)? *Oscillatory entry dynamics across industries (BDS, 1978–2022).* For each of 19 NAICS sectors, the linearly detrended net entry rate exhibits two features outside the AR(1) class: four sectors show significantly negative lag-4 autocorrelation outside the 95% Bartlett (1946) band (incompatible with any scalar AR(1)), and Fisher's g-test rejects the null of no deterministic periodicity at the 5% level in 10 of 19 sectors, with peaks near $T \approx 8$–9 years. The sign-changing autocorrelation pattern is also harder to reconcile with vintage-capital replacement, which produces positively-correlated entry waves rather than the lag-4 sign reversal documented above. Online Appendix D.2 provides the full sector-by-sector results.

Post-shock periodicity in historical impulse responses. I use BEA investment-in-structures data (1930–2024) to test whether one-time shocks trigger oscillations at industry-specific structural frequencies. Manufacturing investment after the Great Depression onset oscillates at $T = 5.0$ years through the early 1970s (Figure 5a; g-test $p = 0.036$ in the 1934–1973 echo window, $p = 0.578$ in 1974–2013), consistent with the bottleneck-friction prediction of decaying post-shock oscillation. Mining/oil investment after the 1979 crisis oscillates at $T = 3.25$ years (Figure 5b; $p = 0.016$). The cross-industry frequency difference is consistent with different (φ, μ) across industries (Online Appendix D.3).

Does the structural separation theorem hold across industries? *Three-moment fingerprint of frequency-welfare separation.* Theorem 3 predicts that the per-industry trio $(\omega_{0,i}, \eta_i, \sigma_{x,i})$ is governed by disjoint primitive subsets ($\omega_{0,i}^2 = \varphi_i/\mu_i$; $\eta_i = \chi_i/\mu_i$; $\sigma_{x,i}^2 = \sigma_{\varepsilon,i}^2/(2\chi_i\varphi_i)$), so cross-industry variation in ω_0 should be uncorrelated with η and negatively correlated with σ_x under approximate μ-uniformity (provided $\sigma_\varepsilon^2/\chi$ does not have offsetting positive covariance with φ). Fitting an AR(2) by Yule-Walker to the band-pass-filtered detrended log level for $N = 15$ NAICS-3 manufacturing industries (Online Appendix D.8), the separation theorem predicts $\rho(\omega_0, \eta) = 0$, while a single-channel persistence model in which one parameter governs both period and damping predicts strong positive (ω_0, η) co-movement. The data give $\rho(\omega_0, \eta) = -0.21$ ($p = 0.46$): the point estimate is consistent with

separation ($\rho = 0$); the test does not have power to reject the single-channel alternative at $N = 15$, but the sign and magnitude run against it. The remaining pairwise correlations are $\rho(\omega_0, \sigma_x) = -0.52$ ($p = 0.05$) and $\rho(\eta, \sigma_x) = -0.16$ ($p = 0.57$); the negative ω_0-σ_x correlation is consistent with μ-uniformity, modulo the offsetting-covariance caveat.

Figure 5: Post-shock periodicity: manufacturing vs. mining/oil investment

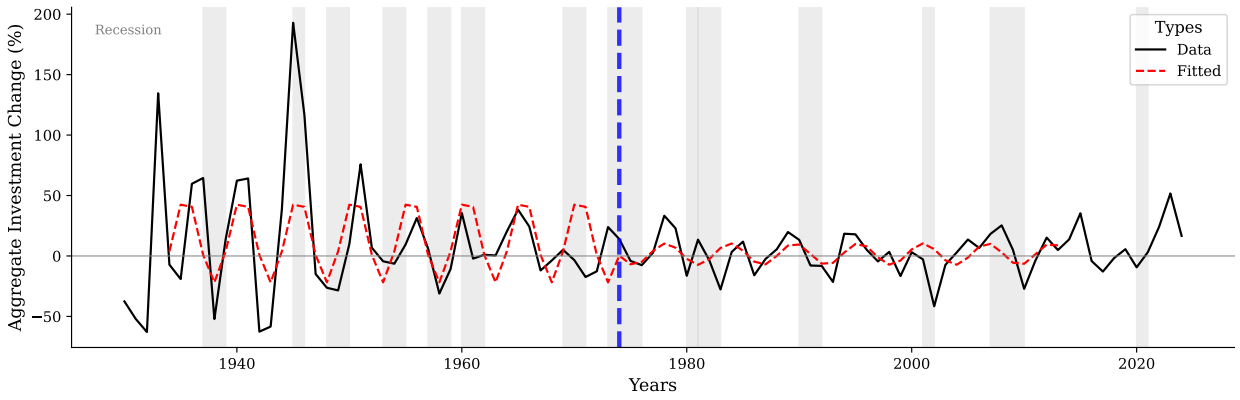

(a) Manufacturing structures

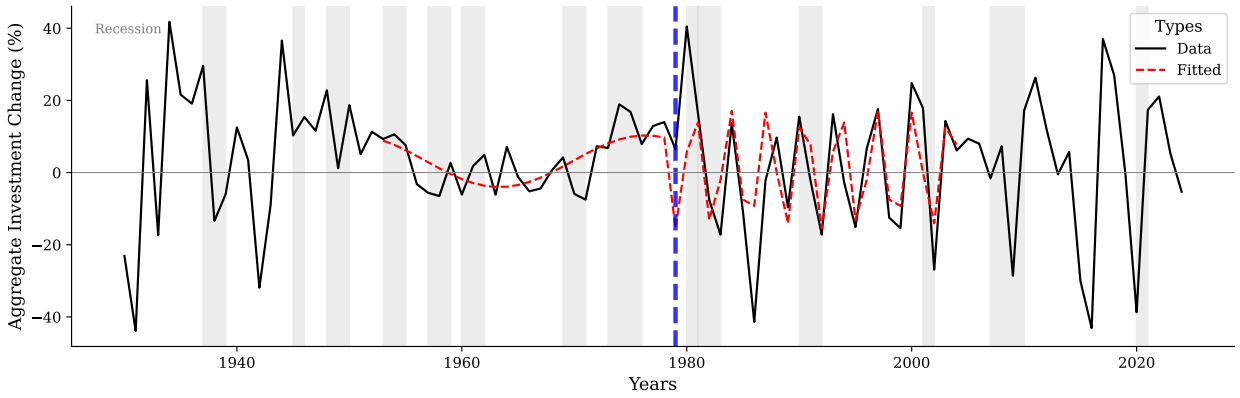

(b) Mining exploration, shafts, and wells

Notes: Annual growth rate of private fixed investment from BEA NIPA Table 5.4.1: panel (a) manufacturing structures (line 16), panel (b) mining exploration, shafts, and wells (line 24). Thick solid line: data. Red dashed line: fitted cosine from Fisher's g-test. Gray bands: NBER recessions. Blue dashed lines mark structural breaks—in (a), the boundary between the post-shock window 1934–1973 and the non-echo window 1974–2013; in (b), the 1979 oil crisis. Manufacturing oscillates at $T = 5.0$ years from the Great Depression onset through the early 1970s ($p = 0.036$); mining/oil oscillates at $T = 3.25$ years post-1979 ($p = 0.016$), consistent with different (φ, μ) across industries.

Do the timing-search predictions about cycle shape and recovery hold? *The post-COVID natural experiment.* Timing search predicts displaced industries overshoot the pre-shock level rather than converge monotonically: the ironing condition forces the

economy past optimal scale before cycling back. After the 2020 shock, 12 of 19 BDS NAICS sectors (63%) overshoot above their pre-COVID linear trend in 2021 or 2022 by at least one pre-COVID standard deviation, against an AR(1)-monotone-convergence baseline of 0.26 (binomial test $p < 0.001$). The baseline is computed by simulating, for each sector, an AR(1) fitted on the linearly-detrended pre-COVID net entry rate (1978–2019) starting from the realized 2020 deviation, and recording the share of 10,000 replications whose 2021–2022 path rises at least one pre-COVID standard deviation above the pre-shock trend purely from innovation noise (Online Appendix D.9 specifies the procedure). The qualitative prediction is borne out at high confidence; the quantitative prediction that overshoot magnitude correlates with concentration is not, with sector-specific COVID exposure dominating cross-sectional variation.

Plucking pattern. Isochrony plus amplitude-proportional recovery velocity (Section 3.5) implies recovery speed scales with recession depth while recovery time does not. For each of seven NBER recessions in FRED INDPRO, 1972–2024, I regress average monthly log growth in the first 12 months post-trough on the peak-to-trough log decline. The slope is +0.046 ($p = 0.029$, $R^2 = 0.65$, Pearson $r = +0.80$). Deeper recessions are followed by faster recoveries, the timing-search analogue of the plucking pattern documented in Friedman (1993); Bordo and Haubrich (2017) and Dupraz et al. (2019).

6 Welfare and industrial policy

6.1 Welfare cost of endogenous cycles

Under the linear-utility household specification of Section 2, household consumption equals aggregate output, so the household welfare loss from cycling is $n_p \, \pi_0 \, \rho^{-1} \, \mathbb{E}[\mathcal{L}]$ to leading order, where $\mathcal{L}$ is the proportional per-firm output deviation. I use "per-firm welfare cost" as shorthand for $\mathcal{L}$ throughout.

The ironing condition provides a natural welfare measure. Along any equilibrium orbit, the equilibrium value level $E = \frac{\varphi}{2} A^2$ is constant by the ironing condition (Equation (12)), and the time-averaged effective output is:

$$\bar{y} = \frac{1}{T} \int_0^T \pi(n(t)) \, \xi(\dot{n}(t)) \, \mathrm{d}t = C = \pi_0 \, e^{-E} = \pi_0 \, \exp\!\left(-\frac{\varphi}{2} A^2\right). \tag{23}$$

Since $A > 0$ in any stochastic equilibrium, $\bar{y} < \pi_0 = \pi(n_p)\xi(0)$: the time-averaged per-firm effective output is strictly below the optimal-scale, zero-adjustment level. The proportional

per-firm effective-output welfare loss from cycling is:

$$\mathcal{L} = 1 - \frac{\bar{y}}{\pi_0} = 1 - e^{-\varphi A^2/2} \approx \frac{\varphi}{2} A^2 \quad \text{for small } A. \tag{24}$$

The welfare cost depends on two structural parameters: φ (the curvature of the profit landscape, determining how costly deviations from optimal scale are) and A (the cycle amplitude, selected in stochastic equilibrium by the shock variance σ_ε^2, the bottleneck friction χ, and the profit-landscape curvature φ). Crucially, $\mathcal{L}$ does not depend on μ. The timing friction determines the *frequency* of cycles but not their welfare cost. Two industries with the same φ and A but different μ cycle at different frequencies yet suffer identical welfare losses.

This gives the Lucas (1987) question a structural answer: where Lucas takes the variance of consumption as a reduced-form object, here $\mathbb{E}[\mathcal{L}] \approx \sigma_\varepsilon^2/(2\chi)$ is a function of measurable primitives (curvature φ from profit landscapes, envelope A from entry-exit cycles), with the welfare cost independent of cycle frequency in the linear small-amplitude limit (Theorem 3; Online Appendix E.1 derives the welfare metric and isolates the supplier-side resource flow $\frac{1}{2}\chi\dot{n}^2$ as a separate accounting object).

Scope of the disjoint-determinants result The exact disjointness $\mathbb{E}[\mathcal{L}] = \sigma_\varepsilon^2/(2\chi)$ holds under linear utility, the per-firm effective-output welfare metric, and exclusion of the supplier's $\frac{1}{2}\chi\dot{n}^2$ resource flow from the cyclical welfare ledger. The same metric choice does load-bearing work earlier in the derivation: it is what makes the equation of motion linear with damping coefficient exactly $\eta = \chi/\mu$ rather than a non-linear function of (φ, μ, χ) (Online Appendix B.2). Both the linear-oscillator structure and the exact disjointness are tied to the same accounting convention. Under CRRA preferences, an aggregate-consumption-equivalent metric, or a welfare measure that internalizes the supplier-side resource cost, the primitives (φ, μ) leak into the welfare formula. Online Appendix E.1 quantifies the leakage as 10–25% at the median across natural alternative metrics: the qualitative ranking—market-structure primitives govern timing, shock variance and bottleneck friction govern welfare—is preserved, but the exact disjointness is not. The most policy-quotable claim of this paper (the CHIPS/IRA counterfactual in Section 6.2) is therefore exact in the tractable benchmark and approximate in the order-of-magnitude robustness range.

Calibrating χ_i for each manufacturing industry yields a mean welfare cost of 0.39% of π_0 across 10 cached industries (range $[0.12\%, 0.70\%]$), comparable to Lucas (1987)- and Reis (2009)-style consumption-equivalent business-cycle welfare costs (Online Appendix E.2). The welfare numbers should be read as central estimates with the same precision caveats the slope test carries: the univariate model-implied second-order fit identifies the forcing

innovation σ_ε/μ, the natural frequency $\omega_0^2 = \varphi/\mu$, and the damping ratio η, but *not* σ_ε separately. Recovering the σ_ε that enters $\mathbb{E}[\mathcal{L}] = \sigma_\varepsilon^2/(2\chi)$ requires the external μ calibration (the Petrongolo–Pissarides matching-elasticity range $\alpha \in [0.3, 0.8]$ used throughout), so the welfare numbers inherit that calibration uncertainty; with ~ 44 annual BDS observations per industry, the per-industry σ_ε is also imprecise. The cross-industry mean of 0.39% is robust to the α range (it shifts by less than 0.10 percentage points across the band per Online Appendix E.2), but the per-industry levels carry both small-sample and calibration uncertainty that should be read symmetrically with the $N = 15$ slope test's wide AR confidence set rather than as point identification.

6.2 Industrial policy and the anatomy of recovery

The frequency-welfare separation (Theorem 3) delivers an instrument-to-outcome map under the per-firm effective-output welfare metric. The structural parameters act on identifiable margins: μ governs timing alone, φ governs timing and severity, χ governs severity and welfare, σ_ε governs severity and welfare, and κ governs only the steady-state ironing level. In the stochastic equilibrium the per-firm welfare cost reduces to $\mathbb{E}[\mathcal{L}] \approx (\varphi/2)\mathbb{E}[A^2] = \sigma_\varepsilon^2/(2\chi)$: both φ and μ cancel, and κ is absent. Table III summarizes the taxonomy.

Table III: Policy anatomy: instrument-to-outcome map

Lever	Trend	Frequency	Severity	Welfare cost	Example
g (productivity)	✓	—	—	—	R&D subsidies
μ (matching friction)	—	✓	—	—	Permitting reform
φ (profit curvature)	—	✓	✓	—	Cluster subsidies
χ (bottleneck friction)	—	—	✓	✓	Capacity expansion
σ_ε (shocks)	—	—	✓	✓	Demand stabilization
κ (sunk cost)	—	—	—	—	Per-firm transfer

Notes: ✓ marks a primitive that moves the column outcome; — marks no effect. Columns are proportional / detrended-log measures: $\omega = \sqrt{\varphi/\mu}$, severity $\sigma_x = \sigma_\varepsilon/\sqrt{2\chi\varphi}$, and *cycle-induced* welfare cost $\mathbb{E}[\mathcal{L}] = \sigma_\varepsilon^2/(2\chi)$ from Theorem 3. The trend column refers to the BGP rate g (Proposition 2); g raises the trend on which welfare is measured—faster growth raises level welfare directly through $z(t) = z_0 e^{gt}$—but does not enter the cyclical welfare cost $\mathbb{E}[\mathcal{L}]$. Welfare-cost effects are stated under the maintained per-firm metric; alternative metrics that include the supplier's resource flow leak μ into the welfare formula at order 10–25% (Online Appendix E.1).

Growth-side and cycle-side levers are structurally separable The taxonomy is a structural decomposition of how primitives map to outcomes, not a one-instrument-one-margin claim about real-world policies. Within the model, the first row of Table III sits orthogonally to the others (Proposition 2): productivity-growth interventions that move g alone leave cycle frequency, severity, and welfare cost untouched, and cycle-side primitives

leave g untouched. The decomposition gives a sharp diagnostic: any change in cycle properties must run through $(\varphi, \mu, \chi, \sigma_\varepsilon)$, and any change in growth must run through g. Real-world policy bundles typically hit several primitives at once—R&D subsidies can alter φ or χ through technology channels, and workforce upgrading can raise g while lowering μ or χ—so the taxonomy is most useful for decomposing a given policy into its margin-by-margin effects rather than for predicting that any specific intervention moves only one outcome.

CHIPS Act counterfactual The CHIPS and Science Act bundles two structurally distinct components in the model: a *growth-side* component (semiconductor R&D incentives, the National Semiconductor Technology Center, workforce-development investments) that raises the BGP rate g, and a *cycle-side* component (manufacturing-grant subsidies that pay for fabrication-plant construction and supplier-side capacity expansion) that lowers χ. The framework classifies these separately. Calibrating $(\varphi, \mu, \chi, \sigma_\varepsilon)$ for NAICS 3344 (Semiconductor and Other Electronic Component Manufacturing) from BDS 1978–2022 yields a baseline spectral period $T = 4.75$ years matching the observed BDS net-entry-rate peak; Table IV reports the model-implied effects of three stylized cycle-side interventions and one combination, all relative to baseline. The growth-side component shifts the trend without entering the table.

Table IV: Counterfactual policy interventions, NAICS 3344

Scenario	Period ratio	Amplitude ratio	Welfare ratio
A. CHIPS-style subsidy: $\chi \times 0.5$	1.00	1.41	2.00
B. Cluster policy: $\varphi \times 1.2$	0.91	0.91	1.00
C. Streamlined permitting: $\mu \times 0.7$	0.84	1.00	1.00
A + B combined	0.91	1.29	2.00

Notes: Calibration and full taxonomy in Online Appendix E.3. Scenario A halves the bottleneck friction χ, a stand-in for an entry subsidy that subsidizes construction and permitting and so weakens the bottleneck-strain mechanism: cycle period unchanged, amplitude rises by $\sqrt{2}$, welfare cost doubles. Scenario B raises φ by 20%: cycle shortens by 9%, amplitude falls by the same factor, welfare unchanged. Scenario C reduces μ by 30%: cycle shortens by 16%, amplitude and welfare unchanged. Combined A+B speeds the cycle and raises both amplitude and welfare cost.

Implications for the current industrial-policy debate The model delivers a counter-intuitive result for the largest class of CHIPS- and IRA-style entry interventions: capacity-expansion subsidies that lower the bottleneck friction χ *raise* per-firm welfare cost by removing the dissipation that bounds cycle amplitude (Scenario A in Table IV). This is exact under the per-firm output metric; under welfare metrics that include the supplier-side resource flow, the welfare cost of capacity expansion is attenuated but the sign is robust (Online Appendix E.1).

The recent industrial-policy turn—the CHIPS and Science Act, the Inflation Reduction Act, the EU Net-Zero Industry Act—has revived the question of which interventions deliver welfare gains and which merely reshape industry dynamics (Juhász et al., 2024; Rodrik, 2004; Bown, 2024); Table III sets out the full structural taxonomy. Productivity-growth interventions (R&D, technology development, workforce upgrading) move g alone, raising the BGP without affecting cycle outcomes—they should be evaluated by growth criteria. Among cycle-side levers, only demand-stabilization measures lower σ_ε and unambiguously reduce welfare cost; permitting reform ($\mu \downarrow$) and cluster support ($\varphi \uparrow$) speed the cycle without moving welfare. This ranking is exact under linear utility and the per-firm output metric; under CRRA preferences or alternative metrics, (φ, μ) leak into welfare at order 10–25% (Online Appendix E.1), so the qualitative ordering holds approximately rather than exactly.

7 Discussion

7.1 Modeling choices and limitations

Five limitations bound the scope of the closed-form result.

(i) Gaussian exactness. The symmetric-cycle prediction of the Gaussian specification is counterfactual; under general hump-shaped profits (Online Appendix B.3), cycles become asymmetric and the frequency formula is accurate to 1.6% at empirically relevant amplitudes (Online Appendix B.5 and C.2).

(ii) Exogenous exit. The model treats exit through the constant-rate δ rather than as a learning- or productivity-driven endogenous decision; combining timing search with Jovanovic (1982)-style learning is a natural extension.

(iii) Linear utility and the per-firm welfare metric. Under CRRA preferences or under welfare metrics that include the supplier's resource flow, (φ, μ) leak into the welfare cost with median magnitude 10–25% (Online Appendix E.1), so the frequency-welfare separation is exact under the maintained metric and approximately exact otherwise.

(iv) Within-manufacturing scope of the empirical identification. The headline test isolates the φ channel under the maintained restriction that μ is approximately uniform across manufacturing 3-digit industries; the substantive treatment is in Section 5.3, with three lines of evidence consistent with the restriction (the recovered $\hat{\beta}_\gamma$ in $[0.50, 0.75]$ within one HC1 standard error of the structural 1/2; capital-side μ proxies returning null coefficients; the microfounded $\mu_i = (1-\alpha)/[\alpha(\delta_i n_{p,i})^2]$ predicting narrow within-manufacturing variation), plus a flagged regulatory-side correlation that does not transmit to the IV slope when added as a control. A direct test of the μ channel outside manufacturing requires industry-friction-

specific proxies at NAICS-3 resolution that current public data do not provide and is left for future data work.

(v) Constant-parameter benchmark vs. time-varying primitives. The closed form $\omega = \sqrt{\varphi/\mu}$ is the *local* natural frequency under constant primitives. Slow time variation in φ or μ—from regulation, finance, supply-chain bottlenecks, technology, or demand composition—generates structurally disciplined but irregular cycles; the empirical prediction is then a ranking statement, that industries with higher persistent φ/μ exhibit higher average dominant frequencies. The cross-sectional slope test in Section 5.3 captures this ranking through the persistent component of γ_i (the 1947 historical IV) and the average spectral peak, identifying the slope provided cross-industry variation in persistent φ/μ dominates transitory variation.

7.2 Extensions

Several extensions merit investigation: combining timing search with Jovanovic (1982)-style learning to endogenize exit; relaxing linear utility to let industries interact through aggregate consumption; endogenizing (φ, μ) through industry investment in linkages or entry infrastructure; combining timing search with financial frictions (Beaudry et al., 2020) to produce models with multiple spectral peaks; and full per-industry GMM/SMM estimation of $(\varphi_i, \mu_i, \chi_i, \sigma_{\varepsilon,i})$ with overidentification testing (the three-moment fingerprint of Section 5.4 is a structural-implication test, not a structural-estimation exercise).

8 Concluding remarks

This paper proposes a structural account of why different industries fluctuate at different frequencies. The central claim is that cycle length is a property of industry market structure—the curvature of the profit landscape and the friction of entry matching—rather than a property of the shock process the industry inherits. Under *timing search*, agents choose *when* to enter a market subject to matching frictions that penalize rapid adjustment; the equilibrium indifference condition across entry dates forces overshoot rather than monotone adjustment toward optimal scale, generating a cycle whose frequency $\omega_0 = \sqrt{\varphi/\mu}$ is pinned by two market-structure primitives alone: the profit-landscape curvature φ and the entry-matching friction μ. The same logic—agents choosing *when* to act under matching frictions—could be extended to other timing decisions such as capital investment, hiring, technology adoption, and financial-market participation, each potentially with its own market-structure-pinned frequency.

U.S. industry data support this prediction. Across 15 manufacturing 3-digit industries,

those with stronger agglomeration cycle at higher frequencies, and the cross-industry slope linking concentration to cycle frequency lands close to the structural value. The pattern is consistent across alternative spectral estimators and survives instrumental-variable identification using historical industrial geography. Standard business-cycle benchmarks can reproduce cross-industry frequency dispersion but do not on their own deliver this slope; the data align with the timing-search account.

Beyond the cross-industry frequency prediction, the model implies that cycle frequency and the welfare cost of cycles are governed by disjoint sets of market primitives. Market structure pins where in the spectrum an industry oscillates, while a separate pair—the variance of profit shocks and the bottleneck friction in the entry pipeline—pins how costly those oscillations are. This separation, which emerges from the same equilibrium condition that pins the frequency, gives a sharp diagnostic for entry-side industrial policy: streamlining entry matching and strengthening agglomeration change cycle speed without moving welfare cost, while only policies that raise bottleneck friction in the entry pipeline or stabilize profit shocks move welfare (Section 6.2). In this view, an industry's cycle frequency depends as much on its market structure as on the shocks it receives.

References

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012). The network origins of aggregate fluctuations. *Econometrica 80*(5), 1977–2016.

Anderson, T. W. and H. Rubin (1949). Estimation of the parameters of a single equation in a complete system of stochastic equations. *Annals of Mathematical Statistics 20*(1), 46–63.

Azariadis, C. (1981). Self-fulfilling prophecies. *Journal of Economic Theory 25*, 380–396.

Bartlett, M. S. (1946). On the theoretical specification and sampling properties of autocorrelated time-series. *Supplement to the Journal of the Royal Statistical Society 8*(1), 27–41.

Beaudry, P., D. Galizia, and F. Portier (2020). Putting the cycle back into business cycle analysis. *American Economic Review 110*(1), 1–47.

Beaudry, P., D. Galizia, and F. Portier (2024). How do strategic complementarity and substitutability shape equilibrium dynamics? NBER Working Paper 32661, National Bureau of Economic Research.

Beaudry, P. and F. Portier (2006). Stock prices, news, and economic fluctuations. *American Economic Review 96*(4), 1293–1307.

Benhabib, J. and R. E. A. Farmer (1994). Indeterminacy and increasing returns. *Journal of Economic Theory 63*, 19–41.

Bordo, M. D. and J. G. Haubrich (2017). Deep recessions, fast recoveries, and financial crises: Evidence from the American record. *Economic Inquiry 55*(1), 527–541.

Bown, C. P. (2024). Modern industrial policy and the WTO. *Annual Review of Economics 16*, 737–765.

Brock, W. A., W. D. Dechert, J. A. Scheinkman, and B. LeBaron (1996). A test for independence based on the correlation dimension. *Econometric Reviews 15*(3), 197–235.

Burdett, K. and D. T. Mortensen (1998). Wage differentials, employer size, and unemployment. *International Economic Review 39*, 257–273.

Cooper, R. W. and J. C. Haltiwanger (2006). On the nature of capital adjustment costs. *Review of Economic Studies 73*, 611–633.

Davis, S. J. and J. Haltiwanger (1992). Gross job creation, gross job destruction, and employment reallocation. *Quarterly Journal of Economics 107*(3), 819–863.

Diamond, P. A. (1982). Efficiency effects of aggregate demand. *Journal of Political Economy 90*(3), 526–553.

Dupraz, S., E. Nakamura, and J. Steinsson (2019). A plucking model of business cycles. NBER Working Paper 26351, National Bureau of Economic Research.

Duranton, G. and D. Puga (2004). Micro-foundations of urban agglomeration economies. In J. V. Henderson and J.-F. Thisse (Eds.), *Handbook of Regional and Urban Economics*, Volume 4. Elsevier.

Eckert, F., K.-l. Lam, A. R. Mian, K. Müller, R. Schwalb, and A. Sufi (2022). The early county business pattern files: 1946–1974. Working Paper 30578, National Bureau of Economic Research.

Ellison, G. and E. L. Glaeser (1997). Geographic concentration in U.S. manufacturing industries: A dartboard approach. *Journal of Political Economy 105*(5), 889–927.

Ellison, G., E. L. Glaeser, and W. R. Kerr (2010). What causes industry agglomeration? evidence from coagglomeration patterns. *American Economic Review 100*(3), 1195–1213.

Farmer, R. E. A. (1999). *The Macroeconomics of Self-Fulfilling Prophecies* (2nd ed.). MIT Press.

Farmer, R. E. A. and J.-T. Guo (1994). Real business cycles and the animal spirits hypothesis. *Journal of Economic Theory 63*(1), 42–72.

Friedman, M. (1993). The "plucking model" of business fluctuations revisited. *Economic Inquiry 31*(2), 171–177.

Gabaix, X. (2011). The granular origins of aggregate fluctuations. *Econometrica 79*(3), 733–772.

Goodwin, R. M. (1967). A growth cycle. In C. H. Feinstein (Ed.), *Socialism, Capitalism and Economic Growth*. Cambridge University Press.

Grandmont, J.-M. (1985). On endogenous competitive business cycles. *Econometrica 53*, 995–1045.

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988). Investment, capacity utilization, and the real business cycle. *American Economic Review 78*(3), 402–417.

Hopenhayn, H. and R. Rogerson (1993). Job turnover and policy evaluation: A general equilibrium analysis. *Journal of Political Economy 101*(5), 915–938.

Hopenhayn, H. A. (1992). Entry, exit, and firm dynamics in long run equilibrium. *Econometrica 60*, 1127–1150.

Hotelling, H. (1931). The economics of exhaustible resources. *Journal of Political Economy 39*(2), 137–175.

Jovanovic, B. (1982). Selection and the evolution of industry. *Econometrica 50*, 649–670.

Jovanovic, B. (2009). Investment options and the business cycle. *Journal of Economic Theory 144*(6), 2247–2265.

Juhász, R., N. J. Lane, and D. Rodrik (2024). The new economics of industrial policy. *Annual Review of Economics 16*, 213–242.

Khan, A. and J. K. Thomas (2008). Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. *Econometrica 76*(2), 395–436.

Kim, S. (1995). Expansion of markets and the geographic distribution of economic activities: The trends in U.S. regional manufacturing structure, 1860–1987. *Quarterly Journal of Economics 110*(4), 881–908.

Krugman, P. (1991). Increasing returns and economic geography. *Journal of Political Economy 99*, 483–499.

Kydland, F. E. and E. C. Prescott (1982). Time to build and aggregate fluctuations. *Econometrica 50*(6), 1345–1370.

Lucas, Robert E., J. (1987). *Models of Business Cycles*. Oxford: Basil Blackwell.

McCall, J. J. (1970). Economics of information and job search. *Quarterly Journal of Economics 84*, 113–126.

Mortensen, D. T. and C. A. Pissarides (1994). Job creation and job destruction in the theory of unemployment. *Review of Economic Studies 61*(3), 397–415.

Petrongolo, B. and C. A. Pissarides (2001). Looking into the black box: A survey of the matching function. *Journal of Economic Literature 39*(2), 390–431.

Pissarides, C. A. (2000). *Equilibrium Unemployment Theory* (2nd ed.). Cambridge, MA: MIT Press.

Reis, R. (2009). The time-series properties of aggregate consumption: Implications for the cost of fluctuations. *Journal of the European Economic Association 7*(4), 722–753.

Rodrik, D. (2004). Industrial policy for the twenty-first century. KSG Working Paper RWP04-047, Harvard Kennedy School.

Rosenthal, S. S. and W. C. Strange (2004). Evidence on the nature and sources of agglomeration economies. In J. V. Henderson and J.-F. Thisse (Eds.), *Handbook of Regional and Urban Economics*, Volume 4. Elsevier.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review 95*(1), 25–49.

Stigler, G. J. (1961). The economics of information. *Journal of Political Economy 69*, 213–225.

Syverson, C. (2004). Market structure and productivity: A concrete example. *Journal of Political Economy 112*, 1181–1222.

Thomson, D. J. (1982). Spectrum estimation and harmonic analysis. *Proceedings of the IEEE 70*(9), 1055–1096.

Wald, A. (1943). Tests of statistical hypotheses concerning several parameters when the number of observations is large. *Transactions of the American Mathematical Society 54*(3), 426–482.

White, H. (1980). A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. *Econometrica 48*(4), 817–838.