Industry Evolution with Cognitive Ladders

Nathan Yang*

April 22, 2026

Abstract

Markets take time to become competitive. I develop a theory in which this reflects cognitive maturation: firms differ in strategic reasoning depth, which advances deterministically with experience. Play converges to Markov perfect equilibrium when strategic feedback is weak, with non-monotone selection along the way. Three results follow. The transition carries independent welfare content, so accelerating cognition (e.g., AI-augmented decision making) can strictly reduce welfare. An interior socially optimal maturation horizon balances this surplus against faster convergence. The maturation process generates a cognitive confounding window during which standard structural estimators cannot separately identify productive and cognitive primitives.

Keywords: Industry Dynamics; Cognitive Hierarchy; Strategic Learning; Bounded Rationality; AI and Competition; Market Structure; Markov Perfect Equilibrium; Oblivious Equilibrium

*University of Illinois Urbana-Champaign, Gies College of Business. Email: ncyang@illinois.edu.

1 Introduction

Markets take time to become competitive. In young industries, firms enter too aggressively, respond sluggishly to rivals, and price in ways that appear insufficiently strategic. When American fast food chains first expanded into Canada in the 1970s, for instance, aggregate entry overshot the eventual steady-state configuration, with waves of openings and closings that only gradually gave way to a stable market structure (Igami and Yang 2016). These patterns are well documented across settings, yet the standard framework for studying industry dynamics assumes that firms play a Markov Perfect Equilibrium at every point in time, as in the canonical model of Ericson and Pakes (1995) and its large-market approximations such as the Oblivious Equilibrium of Weintraub et al. (2008). While the empirical industrial organization literature has made substantial progress in modeling equilibrium outcomes, much less is understood about the transitional processes through which markets approach those outcomes (Pakes 2021). This gap is particularly salient in light of recent evidence documenting rich patterns of firm entry, exit, and reallocation, and a decline in business dynamism over time, both of which highlight the central role of transitional dynamics in shaping market outcomes (Decker et al. 2016).

A natural starting point is that firms differ in their ability to reason strategically. Evidence from a range of contexts shows systematic variation in sophistication: some firms anticipate rivals' responses and adjust accordingly, while others behave in ways closer to myopic or naive decision rules (Goldfarb and Xiao 2011). At the same time, industry dynamics exhibit familiar selection patterns reminiscent of the evolutionary forces emphasized in Jovanovic (1982) and Hopenhayn (1992), as well as in a broader body of work documenting entry, exit, and firm growth over the industry life cycle (Dunne et al. 1988, Klepper and Graddy 1990, Klepper 1996, Bartelsman and Doms 2000). The mechanism proposed here overlaps empirically with these classical learning-and-selection models but operates through a distinct channel. In Jovanovic (1982) and related Bayesian learning frameworks, survivors improve because they revise beliefs about their own fixed type; exit is driven by downward updating on own productivity. Here, survivors revise beliefs about *rivals*, with realized actions (pricing, capacity, entry timing) becoming more strategically calibrated over time. This distinction generates testable implications that a pure learning-about-own-type story does not produce: the U-shaped exit hazard in cognitive state derived below, non-

monotone patterns in entry behavior across cohorts, and markup dynamics tied to the maturation of strategic reasoning rather than to the resolution of productivity uncertainty.

I develop a theory of industry dynamics in which strategic sophistication is a state variable that accumulates with experience. Firms enter with limited ability to reason about competitors and gradually learn to think more deeply. At any point in time, firms best-respond to simplified models of their rivals' behavior, and only as experience accumulates do these models become more refined. Industry dynamics are thus governed by a cognitive ladder: early play reflects shallow reasoning, while later play approaches equilibrium as firms become more sophisticated. The transition toward competitive behavior is not instantaneous; it is itself an equilibrium object shaped by the distribution of strategic ability in the population. Although a firm's cognitive state is not directly observable in administrative data, it can be proxied by firm age relative to market age, managerial tenure or prior-industry experience, repeat-entry status, and adoption of analytics or algorithmic decision tools, linking the theory directly to empirical work on firm lifecycles and organizational sophistication. The framework yields a sequence of results that intuition alone cannot reach. First, the maturation process generates a cognitive confounding window: a structural identification object characterizing the interval during which evolving cognitive dispersion prevents standard estimators from separately identifying productive and cognitive primitives. Second, the exit hazard is U-shaped in cognitive state, and the mechanism is discriminating: shutting down forward-looking behavior eliminates the U-shape entirely, distinguishing the theory from learning-about-own-type accounts that would predict monotone survival. Third, the comparative static on the optimal maturation horizon reverses naive intuition — stronger strategic feedback makes the planner prefer a *longer* horizon, not a shorter one, because the transition surplus persists longer when convergence is slow. These results are not derivable without the model: they depend on the specific interaction of forward-looking behavior, evolving opponent models, and the decay rate of the transition surplus. The framework also delivers a natural application: accelerating cognition (for example through AI adoption) can strictly reduce welfare by compressing the transition surplus, with an interior socially optimal adoption speed that balances this surplus against the informational benefits of faster convergence.

I use the framework to analyze AI adoption as a thought experiment in which the maturation horizon is compressed. Widespread deployment of AI-powered strategic decision tools effectively

accelerates movement along the cognitive ladder, and a growing body of evidence shows that algorithmic agents exhibit deeper opponent modeling and can generate equilibrium outcomes that differ markedly from standard predictions (Brown and MacKay 2023, Deshpande and Jacobson 2026, Rebholz et al. 2025, Fish et al. 2024). Recent work conceptualizes strategy itself as a hierarchy of reasoning tasks in which more sophisticated behavior requires progressively deeper causal and counterfactual reasoning (Chatterji et al. 2026), and advances in multi-agent learning demonstrate that agents capable of reasoning about others' beliefs and learning dynamics can fundamentally alter equilibrium outcomes (Foerster et al. 2017, Yu et al. 2022). AI adoption is thus a natural test case for the framework: it compresses the maturation horizon in a way that would otherwise require exogenous variation across industries.

Two themes of literatures most directly connect with this contribution. Dynamic industrial organization in the Ericson and Pakes (1995) tradition has produced a rich workhorse apparatus, but its models either impose equilibrium play at every point in time or abstract from strategic heterogeneity through simplifying devices such as symmetry restrictions or mean-field approximations (Abbring et al. 2018, Doraszelski and Satterthwaite 2010, Adlakha and Johari 2013, Light and Weintraub 2022, Ifrach and Weintraub 2017). Related work on learning within dynamic oligopoly has emphasized fundamentals such as productivity or costs (Besanko et al. 2010). The static cognitive hierarchy framework of Camerer et al. (2004) treats reasoning depth as a fixed type or endogenous choice rather than a state that evolves with experience. Recent work has begun to extend cognitive hierarchy models to dynamic settings. Lin and Palfrey (2024) develop a dynamic cognitive hierarchy (DCH) solution for extensive-form games in which players update posterior beliefs about opponents' levels after each observed move within a single game, and show that the resulting solution concept violates invariance under strategic equivalence in ways that static CH cannot predict. The DCH framework and the one developed here address different questions with different primitives. DCH tracks the evolution of *beliefs about opponents' levels within a single game* as information accumulates across stages; my framework tracks the evolution of *firms' own reasoning depth across games* as experience accumulates across periods. DCH does not impose the coarse awareness restriction: a DCH player is free to Bayesian-update about the levels of others and to revise play accordingly. My framework imposes a deliberate behavioral restriction on firms' ability to simulate higher levels of their own reasoning, which is what generates the transition dy-

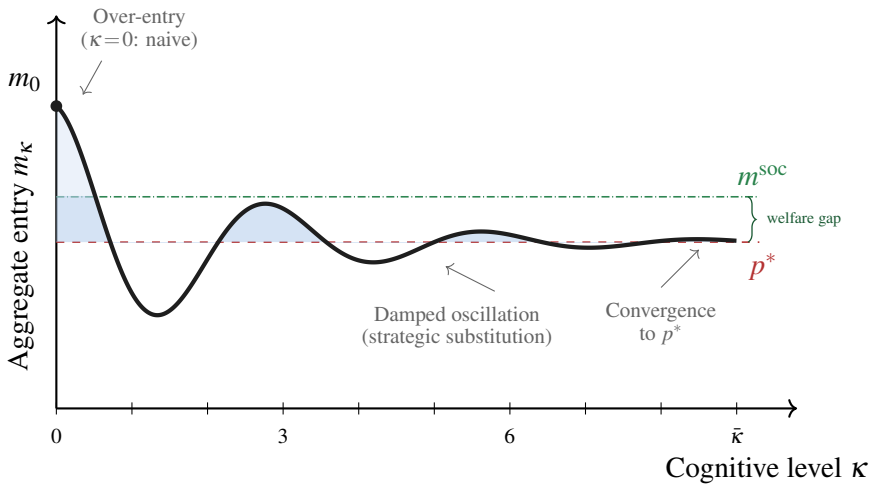
namics. The two frameworks could be combined — firms in a dynamic oligopoly could update beliefs about rivals' cognitive states while their own depth of reasoning matures — but I focus on the second mechanism in isolation, both because it maps cleanly to the industry-dynamics question (how does aggregate play change as a cohort ages?) and because shutting down within-game belief updating sharpens the comparative statics around the transition horizon. My framework sits at the intersection of two traditions: I embed a cognitive hierarchy in a dynamic environment with an evolving population state, connecting naturally to recent work on distributional equilibria in large dynamic games in which the evolution of the population state is itself the central object (Balbus et al. 2022). Here the distribution that evolves is not only over productive states but also over cognitive states, linking strategic reasoning directly to aggregate dynamics. A closely related stream, focused on identification rather than modeling, studies dynamic games when players' beliefs are not in equilibrium and develops nonparametric tools for jointly identifying payoffs and beliefs without imposing equilibrium restrictions (Aguirregabiria and Magesan 2020). My framework complements that program. In particular, I impose a specific structural form on belief bias and trade the generality of their identification approach for sharper predictions about the path of bias, its welfare consequences, and its implications for structural estimation.

2 The Model

Before turning to the formal framework, it is useful to preview the central mechanism. Figure 1 illustrates how aggregate entry evolves as the population of firms advances along the cognitive ladder. Naive firms (i.e., at cognitive level $\kappa = 0$) over-enter relative to both the MPE entry probability p^* and the social optimum m^{soc}, generating an initial burst of excess competition. As firms gain experience and their strategic reasoning deepens, aggregate entry converges toward p^* with damped oscillations. This pattern reflects strategic substitution: each successive cohort partially corrects the mistakes of its predecessors but does so imperfectly, leading to smaller errors that alternate in direction.

The shaded region between p^* and m^{soc} represents the transition surplus. This surplus arises because cognitive heterogeneity temporarily moves aggregate behavior closer to the social optimum than the MPE would deliver in a fully mature industry. In particular, the initial overshoot

Figure 1: Aggregate Entry Dynamics Along the Cognitive Ladder



Notes: At $\kappa = 0$, naive firms over-enter relative to both the MPE entry probability p^* and the social optimum m^{soc}. As cognitive depth increases, aggregate entry oscillates around p^* with geometrically decaying amplitude, driven by strategic substitution under $\rho^* < 1$. The shaded region between p^* and m^{soc} represents the transition surplus $S(\kappa) = W^{\text{CLE}}(\kappa) - W^{\text{MPE}}$, corresponding to periods in which cognitively heterogeneous play generates higher welfare than the fully strategic equilibrium. The figure depicts the natural maturation path; under AI acceleration, the same oscillation is compressed into a shorter horizon $T_{AI} < T_{\text{nat}}$, eliminating the portion of the transition surplus that would otherwise accrue over $[T_{AI}, T_{\text{nat}}]$.

in entry brings competition nearer to m^{soc}, generating welfare gains that disappear once all firms behave strategically. Accelerating cognitive maturation (i.e., compressing the transition), for example through the adoption of AI-based decision tools, eliminates part of this surplus by truncating the path of adjustment and moving the industry more quickly to p^*.

The model developed below formalizes these components: the deterministic progression along the cognitive ladder (Section 2.2), the equilibrium concept governing behavior at each level (Section 2.3), and the resulting implications for convergence, selection, and welfare (Section 3).

2.1 Environment

Time is discrete, $t = 0, 1, 2, \ldots$. The industry evolves through the interaction of a finite set of active firms $\mathcal{N}_t$, where the number of firms $N_t = |\mathcal{N}_t|$ is itself endogenous through entry and exit. The key departure from the standard framework is that firms are heterogeneous not only in their productive characteristics, but also in their ability to reason strategically.

Each firm $i \in \mathcal{N}_t$ is therefore described by two state variables. The first is a productive state $x_{it} \in \mathcal{X} \subset \mathbb{R}_+$, which captures conventional sources of heterogeneity such as productivity, cost efficiency, or product quality. This state evolves with investment and determines the firm's payoff-

relevant fundamentals, as in standard models of industry dynamics. The second is a cognitive state $\kappa_{it} \in \{0, 1, \ldots, \bar{\kappa}\}$, which represents the firm's current depth of strategic reasoning. A higher κ corresponds to a greater ability to anticipate and respond to rivals' behavior. The parameter $\bar{\kappa} < \infty$ imposes a maximal level of sophistication, reflecting a bound on strategic capacity even in the long run.

The industry state at time t is given by the joint distribution of these characteristics across firms,

$$s_t = \{(x_{it}, \kappa_{it})\}_{i \in \mathcal{N}_t} \in \mathcal{S}.$$

This object summarizes both the economic fundamentals of the industry and the distribution of strategic sophistication. In contrast to standard models, where the state tracks only productive characteristics, here the evolution of the cognitive distribution is an independent driver of industry dynamics.

Firms choose actions $a_{it} \in \mathcal{A}$, where $\mathcal{A}$ is a compact metric space. Per-period payoffs are given by $\pi(x_{it}, \hat{a}_{-i,t}, a_{it})$, where $\hat{a}_{-i,t}$ denotes the distribution of opponents' actions. This formulation makes explicit that payoffs depend on rivals' actions rather than directly on their cognitive states; the cognitive composition of the industry affects outcomes indirectly through the actions it induces. Payoffs are symmetric across firms and continuous in $(x, \hat{a}_{-i}, a)$.

2.2 State Transitions

The two state variables evolve according to distinct mechanisms, reflecting the separation between economic fundamentals and strategic ability.

The productive state evolves as in Ericson and Pakes (1995). Given current state x_{it} and action a_{it}, the next-period productive state is drawn from a transition kernel:

$$x_{i,t+1} \sim Q_x(\cdot \mid x_{it}, a_{it}).$$

Investment stochastically upgrades the state, so firms can influence their productive capabilities through their actions. The kernel Q_x satisfies standard regularity conditions (i.e., continuity in a and the Feller property), ensuring that the dynamic program is well behaved. This component

of the model is conventional and captures the usual sources of persistence and heterogeneity in industry dynamics.

In contrast, the cognitive state evolves exogenously and deterministically according to

$$\kappa_{i,t+1} = \min(\kappa_{it} + 1, \bar{\kappa}). \tag{1}$$

Entry cognitive states $\kappa_{i,0}$ are drawn from an initial distribution F_0 supported on $\{0, 1, \ldots, \bar{\kappa}\}$. This transition rule reflects the central modeling choice: reasoning depth is not a decision variable but an accumulation process. Firms do not choose how deeply to reason in each period; instead, strategic ability improves passively with experience at a fixed rate. As a result, all firms move up the cognitive ladder over time, but they do so from potentially different starting points.

Exit decision and idiosyncratic profitability shock. At the end of each period, before the next period's state is realized, firm i receives an idiosyncratic profitability shock $\varepsilon_{it} \sim F_\varepsilon$ (i.e., i.i.d. with continuous CDF, mean zero, and independent of (x_{it}, κ_{it})), capturing unobserved cost realizations, demand fluctuations, or operational disruptions that are orthogonal to the firm's state.

The firm exits if and only if its total realized continuation value falls below the scrap value $\underline{V}$:

$$\text{exit} \iff V^{\text{true}}(x_{it}, \kappa_{it}) + \varepsilon_{it} < \underline{V}.$$

This formulation makes exit stochastic rather than deterministic. The conditional exit hazard at productive state x and cognitive state κ is

$$h(\kappa \mid x) \equiv F_\varepsilon(\underline{V} - V^{\text{true}}(x, \kappa)),$$

which is a continuous, strictly monotone function of $V^{\text{true}}(x, \kappa)$. Intuitively, firms with higher continuation values are less likely to exit, but even identical firms may make different exit decisions due to idiosyncratic shocks. This feature ensures that exit behavior varies smoothly across firms (i.e., it is not a degenerate threshold rule), which aligns with empirical exit hazards observed in the data.

Because exit is stochastic, there is no single value of x that deterministically separates exiters

from survivors. References to "x above the exit threshold" in subsequent assumptions should be interpreted as shorthand for x in the region where $V^{\text{true}}(x,\kappa) > \underline{V}$ with positive probability: formally, the set $\{x : V^{\text{true}}(x,\kappa) > \underline{V} - \bar{\varepsilon}\}$ for some $\bar{\varepsilon}$ in the support of F_ε.

The coarse awareness restriction. The transition $\kappa' = \kappa + 1$ is deterministic and known to the firm. A natural concern is that a forward-looking firm might anticipate this transition and effectively "skip ahead" by computing the implications of its future, more sophisticated self. If this were possible, the cognitive ladder would collapse back to rational play. The model rules out this possibility by imposing a restriction on what firms are able to represent and evaluate.

Definition 1 (Truncated Belief Operator)**.** For a firm at cognitive level κ, the belief operator $\mathscr{B}^\kappa$ maps the history of observables into a belief over opponents' strategies. The truncation condition restricts $\mathscr{B}^\kappa$ to the class of belief-forming processes that can be expressed as depth-κ iterated best responses applied to σ_0. Formally, $\mathscr{B}^\kappa$ has range contained in $\{\sigma^{(0)}, \sigma^{(1)}, \ldots, \sigma^{(\kappa)}\}$: the firm can represent strategies up to depth κ but cannot construct or evaluate $\sigma^{(\kappa+1)}$ without first advancing its cognitive state.

This restriction applies recursively in the continuation problem. When the firm solves its Bellman equation, it evaluates continuation values using $\mathscr{B}^\kappa$-consistent beliefs. In particular, the continuation value $V^{(\kappa)}$ reflects payoffs under the belief that future opponents play $\sigma^{(\kappa-1)}$ at all future nodes, because the belief operator does not update with the anticipated increase in the firm's own cognitive state. Although the firm knows that its reasoning ability will improve in the next period, it cannot revise its model of opponents' behavior to reflect what a more sophisticated self would infer. The firm's information set at time t is therefore effectively partitioned by κ_{it}: objects that require level-$(\kappa+1)$ reasoning are not accessible until the firm reaches that level. This restriction prevents forward induction from unraveling the model and ensures that each level's optimization problem is well defined.

A direct implication is that the level-κ firm is systematically mistaken about its own future payoffs as well as about rivals. When the firm transitions to level $\kappa+1$, its realized continuation value will be $V^{(\kappa+1)}$, not $V^{(\kappa)}$. By using $V^{(\kappa)}$ in the current Bellman equation, the firm operates under a projection bias about its future strategic capacity (Loewenstein et al. 2003): it correctly anticipates that its reasoning will improve, but evaluates that improvement using its current, less sophisticated

perspective. In this sense, the firm is naive about its own future cognitive limitations (i.e., it does not anticipate that its future self will solve a different problem with different beliefs), consistent with the notion of naivety in O'Donoghue and Rabin (1999). This is a deliberate modeling choice rather than an inconsistency, and it is the dynamic analogue of the standard assumption in static cognitive hierarchy models that agents do not anticipate higher-order reasoning by others.

Status of the deterministic ladder: a restricted hypothesis class. The combination of deterministic advancement and the truncated belief operator is a reduced-form behavioral restriction, not a microfoundation. A common objection is that the CLE represents a failure of optimization: why doesn't a forward-looking firm simulate higher-level reasoning and collapse the ladder to full sophistication? The more accurate reading is that the CLE is not a failure of optimization but a restriction on the space of strategies a firm can represent at a given level of experience. To make this precise, define the *level-κ strategy class* $\mathscr{S}^{(\kappa)}$ recursively: $\mathscr{S}^{(0)}$ is the set of uniform-random policies, and $\mathscr{S}^{(\kappa)}$ is the set of policies obtainable by applying the single-agent best-response operator to any belief concentrated on $\mathscr{S}^{(\kappa-1)}$. The CLE policy $\sigma^{(\kappa)}$ is the unique fixed point of the best-response operator restricted to $\mathscr{S}^{(\kappa)}$. The restriction is closure under κ-step best-response composition, not an assumption about cognitive effort or computational cost. Importantly, $\mathscr{S}^{(\kappa)}$ is not an ad hoc restriction imposed from outside the model: it is generated endogenously by iterated application of the best-response operator starting from $\mathscr{S}^{(0)}$, and the level-κ policy is the rational response to the previous layer of the resulting hierarchy. What the model *does not* impose is bounded foresight, mistaken beliefs in the Bayesian sense, or suboptimality within $\mathscr{S}^{(\kappa)}$; firms optimize perfectly within the class, and the class itself is the primitive. This distinguishes the CLE sharply from three neighboring constructs:

- *Bounded rationality as computational cost* (Rubinstein 1998) treats reasoning as a costly process and compares strategies on a cost-benefit basis. The CLE imposes no cost; the restriction is on representability.

- *Incorrect beliefs* (Aguirregabiria and Magesan 2020) treat firms as optimizers under a prior that may be inconsistent with the true environment. The CLE can be written in this language only if the prior is further restricted to be degenerate on $\mathscr{S}^{(\kappa-1)}$, which is a much stronger requirement than non-equilibrium beliefs alone.

- *Limited foresight* models firms that discount beyond the standard rate or truncate the horizon. The CLE preserves the full horizon and the standard discount rate (see "Long-horizon but naive" below); what is truncated is the class of strategies firms can represent, not the horizon over which they project.

The following proposition makes the restricted-hypothesis-class framing formal.

Proposition 1 (CLE as restricted optimization)**.** *For each* $\kappa \in \{0, 1, \ldots, \bar{\kappa}\}$*, the CLE policy* $\sigma^{(\kappa)}$ *is the unique solution to the constrained Bellman problem*

$$\sigma^{(\kappa)} = \arg\max_{\sigma \in \mathscr{S}^{(\kappa)}} \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \pi(x_t, m_t^{(\kappa-1)}) \,\middle|\, \sigma\right],$$

where $m_t^{(\kappa-1)}$ *is the aggregate play induced when all rivals use* $\sigma^{(\kappa-1)}$ *and the expectation is taken under the true transition kernel* Q_x*. Equivalently,* $\sigma^{(\kappa)}$ *is the unique fixed point of the unconstrained Bellman operator applied to beliefs concentrated on* $\sigma^{(\kappa-1)}$*.*

The proof is in Appendix 4. The content of Proposition 1 is that the CLE policy is the full-optimization solution within a restricted strategy class. Two robustness extensions preserve this interpretation. First, if cognitive advancement is stochastic (e.g., with probability λ_κ per period of advancing one rung), the convergence and welfare results survive in expectation, with $\bar{\kappa}$ replaced by the expected cap. Second, if firms can invest in cognitive advancement at a cost, the model admits an optimal-investment extension in which the planner's problem of Proposition 9 becomes a dual to firms' private investment decisions. Neither extension changes the fundamental object: a family of restricted strategy classes indexed by experience, with full optimization within each class.

Long-horizon but naive. A low-κ firm is not short-sighted. It discounts the future at the standard rate β and builds plans that extend over the full horizon. What distinguishes it is the model under which those plans are formed: long-range projections rest on a belief that rivals behave according to $\sigma^{(\kappa-1)}$ at every future node, and this belief does not update even as the firm anticipates its own cognitive advancement. A startup founder building a five-year financial projection illustrates the distinction. The founder discounts future cash flows normally and plans carefully over the full

horizon. The projections, however, rest on competitive assumptions ("rivals will respond slowly," "the incumbent will not re-price") that are themselves products of shallow strategic reasoning. As the founder accumulates experience and advances along the cognitive ladder, the projections are revised and the earlier plan is revealed to have been long-horizon but naive, not short-sighted.

2.3 Strategic Play and the Cognitive Ladder Equilibrium

A firm with cognitive state κ behaves according to a level-κ decision rule, where higher κ corresponds to a deeper ability to anticipate and respond to rivals.

Level-0. Level-0 firms follow a fixed exogenous policy $\sigma_0 \in \Delta(\mathscr{A})$, interpreted as a myopic benchmark (i.e., a rule that ignores strategic interaction). Concretely, the firm behaves as if no rivals are present, so the binary entry decision is $v - c(1+0) + \varepsilon_i > 0$, yielding a stand-alone payoff component of $v - c$. In the binary entry game (see Section 3.2 for the full specification with N firms), this gives $\sigma_0 = \Lambda(\lambda(v-c))$, the entry probability under zero competitive crowding from rivals. Since $v > c > 0$ and Λ is strictly increasing, this implies $\Lambda(\lambda(v-c)) > \Lambda(\lambda(v-c-c(N-1)p^*)) = p^*$, so level-0 behavior generates over-entry relative to equilibrium.

Level $\kappa \geq 1$. A level-κ firm behaves strategically, but only up to its current reasoning depth. It best-responds to the belief that all opponents are one step less sophisticated (i.e., at level $\kappa - 1$):

$$a_{it}^{(\kappa)} \in \arg\max_{a \in \mathscr{A}} \left\{ \pi(x_{it}, \hat{a}_{-i,t}^{(\kappa-1)}, a) + \beta \, \mathbb{E}\left[V^{(\kappa)}(x_{i,t+1}) \mid x_{it}, a \right] \right\}, \tag{2}$$

where $\hat{a}_{-i,t}^{(\kappa-1)}$ denotes the distribution of opponents' actions implied by level-$(\kappa-1)$ behavior. The continuation value $V^{(\kappa)}(x)$ is evaluated under level-κ beliefs, and under the coarse awareness restriction the value function takes x alone as its argument: the firm correctly knows its cognitive state will advance, but evaluates the resulting payoffs using its current level-κ perspective without substituting the level-$(\kappa+1)$ value function. The expectation integrates over the known productive-state transition Q_x.

The belief structure embeds a deliberate restriction: a level-κ firm holds a degenerate belief that all opponents behave as level-$(\kappa-1)$ types. This differs from formulations in which firms

assign probability mass across multiple levels. The restriction is structural rather than cosmetic. Representing a distribution over opponent types would require reasoning about why others might be at different levels, which is itself a higher-order task that lies beyond level-κ cognition. As a result, the firm's perception of the industry is systematically simplified. Several of the paper's key results, including the welfare surplus and the non-monotone selection pattern, rely on this misperception of heterogeneity rather than on bounded depth alone. Allowing richer belief distributions would attenuate these effects by partially restoring awareness of the true environment.

Definition 2 (Cognitive Ladder Equilibrium)**.** A *Cognitive Ladder Equilibrium* (CLE) consists of value functions $\{V^{(\kappa)}\}_{\kappa=0}^{\bar{\kappa}}$, strategy functions $\{\sigma^{(\kappa)}\}_{\kappa=0}^{\bar{\kappa}}$, and an industry state process $\{s_t\}_{t\geq 0}$ such that: (i) $\sigma^{(0)} = \sigma_0$ is the exogenous level-0 policy; (ii) for each $\kappa \geq 1$, $\sigma^{(\kappa)}$ solves the level-κ Bellman equation (2) given level-$(\kappa-1)$ opponent behavior; (iii) exit is individually optimal given the true continuation value $V^{\text{true}}(x_{it}, \kappa_{it})$ (as described in Section 2.2), while action choices are individually optimal given the cognitive-state-specific subjective value functions $V^{(\kappa)}$; and (iv) the industry state evolves consistently with individual behavior and the transition rules for (x, κ).

The CLE is a behavioral equilibrium in the sense that firms optimize given a misspecified model of the environment, where the misspecification arises from cognitive limitations rather than informational frictions. This connects the framework to the broader behavioral IO literature on naive agents, while extending it to a dynamic setting in which the degree of misspecification evolves over time.

Modeling boundary: no meta-learning about cognition. The CLE assumes that firms do not infer rivals' cognitive levels from observed behavior and do not update beliefs about the distribution of sophistication in the population. In particular, a level-κ firm does not attempt to imitate more sophisticated competitors even if their behavior appears systematically different. Allowing such inference or imitation would effectively accelerate the cognitive ladder, since firms could move toward equilibrium behavior without waiting for the deterministic progression of κ. Ruling out this channel isolates the role of cognitive maturation itself and keeps the transition dynamics transparent.

Common knowledge in the CLE. The CLE retains common knowledge of the game structure (the payoff function, the action space, and the transition kernel Q_x) and of the cognitive ladder itself (the recursion $\sigma^{(\kappa)} = \text{BR}(\sigma^{(\kappa-1)})$). It departs from the standard MPE assumption of common knowledge of rationality: a level-κ firm does not believe its rivals are fully rational but rather one step less sophisticated. It also departs from common knowledge of the cognitive distribution: the cross-sectional frequency $\mu_t(\kappa)$ is an object known to the analyst but not to firms, which treat opponents as a homogeneous level-$(\kappa-1)$ population. These epistemic weakenings are what allow the CLE to generate transition dynamics distinct from the MPE while retaining a well-defined best-response structure at each level.

Subjective versus objective behavior. A useful distinction is between the behavior firms expect and the behavior that actually occurs. A level-κ firm believes that opponents behave according to level-$(\kappa-1)$ rules, which determines its subjective environment. Actual outcomes, however, reflect the full distribution of cognitive states in the population. The gap between perceived and realized behavior represents the firm's cognitive error. This error shrinks as κ increases, which is the fundamental force driving convergence in the model.

Existence via forward induction. The CLE exists for any finite $\bar{\kappa}$ and is constructed sequentially. Level-0 behavior is given exogenously. Given level-$(\kappa-1)$ behavior, the level-κ firm solves a standard dynamic program against a known environment, yielding a unique solution under standard regularity conditions. Iterating this construction yields the full equilibrium. The key intuition is that each level solves a well-defined problem against a fixed lower-level policy, which avoids the circularity that typically complicates dynamic games.

The CLE nests the standard MPE as a limiting case (i.e., when all firms are sufficiently sophisticated). If all firms begin at a high enough level, behavior is effectively indistinguishable from equilibrium play. When cognitive capacity is bounded below this threshold, even a mature industry does not reach full equilibrium, and persistent deviations remain.

Robustness to the level-0 specification. The role of σ_0 is to anchor the initial condition of the cognitive ladder. The qualitative results on convergence, selection, and path dependence depend on the structure of the best-response mapping rather than the specific choice of σ_0. In contrast, welfare

comparisons depend on whether level-0 behavior over- or under-reacts relative to equilibrium. The maintained specification focuses on over-entry (i.e., naive optimism), which is empirically common, but the framework accommodates alternative anchors with corresponding changes in welfare implications.

Lemma 1 (Transition Suppression is Robust to σ_0)**.** *If σ_0 generates aggregate entry $m_0 \in (m^{mon}, m^{comp})$ with m_0 and p^* on opposite sides of m^{soc} (i.e., initial entry exceeds equilibrium entry while the social optimum lies between them), then there exist periods during the transition phase in which $W^{CLE}(t) > W^{MPE}$.*

The analysis also accommodates the case in which level-0 firms under-enter (i.e., pessimism about profitability). Under strategic substitution, an initial period of under-entry is followed by over-entry in the next period, and the sequence oscillates around p^* with declining amplitude. When $p^* < m^{\text{soc}}$, these oscillations will intermittently place aggregate entry above m^{soc}, generating transient welfare surpluses that acceleration would eliminate. Consequently, the welfare effect of acceleration when $m_0 < p^*$ is not unambiguously positive; it depends on the relative magnitudes of the surpluses lost and retained along the oscillating path. The baseline maintains the over-entry specification because it aligns with the most common empirical patterns in early industry dynamics and produces the cleanest characterization of the transition surplus.

3 Key Insights

3.1 What the CLE Framework Delivers that DCH Cannot

Before turning to the headline results, I establish formally what the CLE framework delivers that the dynamic cognitive hierarchy of Lin and Palfrey (2024) cannot. The three central objects of this paper — the transition surplus $S(t)$, the U-shaped selection hazard, and the identification distortion Ψ_t — are all functionals of the time evolution of the aggregate cognitive distribution μ_t. The following lemma shows that in DCH, this evolution is identically zero.

Lemma 2 (Aggregate Degeneracy of DCH)**.** *Consider the DCH framework of Lin and Palfrey (2024) applied to a population of N firms, each a fixed-type level-κ reasoner drawn from a popu-*

lation distribution μ. Then the aggregate cognitive distribution satisfies $\mu_t = \mu$ for all t. Consequently, the three central objects of this paper are degenerate in DCH:

(i) Transition surplus: $S(t) \equiv W^{CLE}(t) - W^{MPE} = W^{DCH}(\mu) - W^{MPE}$ is time-invariant, eliminating the welfare dynamics that Proposition 6 characterizes.

(ii) U-shaped hazard: The exit hazard $h(\kappa \mid x)$ in Proposition 5 is a mapping from fixed type κ to hazard level, not a path traced by firms as they age. The non-monotonic selection pattern is not defined in DCH because cognitive states do not evolve within firms.

(iii) Identification distortion: The cognitive confounding window $\mathscr{T}^{\eta}$ in Proposition 10 collapses to the empty set or the full positive real line, depending on whether $\Psi_t < \eta$ holds for the (time-invariant) distribution μ. Cohort-age-dependent identification variation does not exist.

The proof is in Appendix 4 and consists of verifying that DCH's within-game belief updating operates on *beliefs about opponents' levels* with fixed own-levels, so the population distribution of own-levels is a stationary object. The intuition is direct: DCH adds sophistication to within-game play but does not supply a law of motion for the cognitive composition of the population. The CLE framework is the complement. It specifies an explicit cross-period law of motion for μ_t, making the three headline objects well-defined functionals of that law.

This separation clarifies where the two frameworks substitute and where they complement. On the *within-game* dimension (beliefs about rival sophistication), DCH delivers sharper predictions than the CLE's coarse-awareness restriction. On the *across-period* dimension (evolution of own cognitive state), the CLE generates dynamics that DCH cannot reach by construction. The combined framework — DCH-style within-game belief updating layered on CLE-style across-period cognitive maturation — is a natural next step but falls outside the scope of the current paper.

3.2 Convergence: To MPE, to OE, or Neither

The central question of this section is how cognitive maturation translates into aggregate play and whether that play converges to equilibrium. The mechanism is easiest to see in a canonical binary

entry game, which also furnishes the leading example for the rest of the paper. The general-case extension is recorded in a remark at the end of the subsection.

A canonical binary entry game. Consider an industry with N firms. Firm i enters if

$$v - c \cdot (1 + a_{-i}) + \varepsilon_i > 0,$$

where v is stand-alone value, c captures competitive business-stealing effects, a_{-i} is the number of entering rivals, and ε_i is a logistic shock with precision λ. The MPE entry probability p^* is the fixed point of the stationary Bellman best-response operator:

$$p^* = \Lambda(\lambda(v - c - c(N-1)p^*)), \qquad \Lambda(z) = \frac{1}{1+e^{-z}}, \tag{3}$$

which admits a unique solution in $(0,1)$ under standard conditions.

Level-κ play in this setting is a scalar entry probability $p_\kappa \in (0,1)$ generated by the recursion $p_\kappa = \Lambda(\lambda(v - c - c(N-1)p_{\kappa-1}))$, with $p_0 = \sigma_0$ the exogenous anchor. The derivative of this best-response mapping evaluated at p^* is

$$J^* = -\lambda c(N-1)p^*(1-p^*), \qquad \rho^* = |J^*| = \lambda c(N-1)p^*(1-p^*).$$

The scalar ρ^* measures how strongly a small error in beliefs propagates through one step of best-responding. When $\rho^* < 1$, errors dampen at each level and the sequence $\{p_\kappa\}$ converges geometrically to p^*. When $\rho^* > 1$, the interior fixed point is unstable. The sign of J^* determines whether convergence is oscillatory ($J^* < 0$, strategic substitution) or monotone ($J^* > 0$, strategic complementarity). Crucially, ρ^* is determined entirely by payoff primitives, so the speed and stability of convergence are features of the economic environment rather than of the cognitive process itself.

Assumption 1 (Regularity)**.** The per-period payoff function $\pi(x, \hat{a}_{-i}, a)$ is twice continuously differentiable in (x,a) and continuous in $\hat{a}_{-i}$; the state space $\mathscr{X}$ is compact; and the transition kernel Q_x satisfies a uniform Lipschitz condition in a. Interior solutions to the level-κ Bellman equation exist for all κ.

Assumption 2 (Fixed Cohort)**.** For the purpose of Proposition 2, the industry is composed of a fixed cohort of N firms that enter at $t = 0$ with cognitive states $\{\kappa_{i,0}\}_{i=1}^{N}$ drawn from F_0, and no new entry occurs after $t = 0$. Each firm ages deterministically so that $\kappa_{it} = \min(\kappa_{i,0} + t, \bar{\kappa})$.

Proposition 2 is established under Assumption 2, which halts new entry and produces a clean cohort-aging structure. Under this assumption, the cognitive distribution shifts upward rigidly over time, and the opponents' environment can be treated as an exogenous (if time-varying) process from the perspective of each level-κ firm. With ongoing entry, new firms continually arrive at low cognitive levels, preventing the distribution from converging to a point mass at $\bar{\kappa}$. In that case, the industry approaches a stationary distribution with persistent cognitive dispersion, reflecting a balance between inflow at low levels and maturation toward $\bar{\kappa}$. Formalizing this stationary regime is a natural extension; the core convergence mechanism is the same, but it operates in steady state rather than in finite time.

Proposition 2 (Convergence to MPE)**.** *Under Assumptions 1 and 2 and the contraction condition $\rho^* < 1$, there exists a finite saturation depth K^* such that for any $\bar{\kappa} \geq K^*$ and any initial level-0 anchor σ_0 within the basin of attraction of the MPE fixed point p^*, the time-t aggregate distribution of play under the CLE satisfies*

$$\|m_t - m^{MPE}\| \leq C(\rho^*)^t + \|\sigma^{(\bar{\kappa})} - \sigma^*\|$$

for some $C < \infty$, where the residual term $\|\sigma^{(\bar{\kappa})} - \sigma^*\|$ represents the irreducible gap between the capped level-$\bar{\kappa}$ policy and the MPE. For any tolerance $\varepsilon > 0$, choosing $\bar{\kappa} \geq K^*(\varepsilon)$ ensures $\|\sigma^{(\bar{\kappa})} - \sigma^*\| \leq \varepsilon$, so that for practical purposes the bound reduces to $\|m_t - m^{MPE}\| \leq C(\rho^*)^t + \varepsilon$.

This proposition illustrates that when strategic feedback is sufficiently weak, each step up the cognitive ladder reduces the discrepancy between perceived and actual behavior. As firms become more sophisticated, their actions become closer to equilibrium best responses, and the aggregate outcome converges up to the residual error induced by the finite cap. The rate of convergence is geometric and governed by ρ^*, so environments with stronger competitive interactions converge more slowly.

Proposition 3 (OE Validity Requires Cognitive Maturity)**.** *Let $\underline{\kappa}_t$ denote the minimum cognitive state among active firms at time t. In the large-market limit $N \to \infty$, the standard stationary OE approximation is valid only when (i) the light-tail condition on the stationary distribution of x holds, and (ii) the cognitive state distribution has effectively saturated (i.e., $\underline{\kappa}_t \geq K^*$).*

When condition (ii) fails, the industry state distribution is non-stationary due to ongoing cognitive advancement, and the OE approximation mischaracterizes transitional dynamics. In this case, a formulation that tracks the evolution of cognition is required. The intuition is that the OE approximation relies on stationarity of the environment faced by firms. When cognition is still evolving, the environment itself is drifting over time, so averaging across states obscures the transition dynamics.

Proposition 4 (Divergence from Interior Equilibrium under Complementarity)**.** *Under strategic complementarity (so $J^* > 0$), the contraction condition $\rho^* < 1$ may fail. When $\rho^* > 1$, the interior MPE p^* is unstable: for any $p_0 \neq p^*$, the sequence $p_\kappa = BR^{\circ\kappa}(p_0)$ converges monotonically to one of the boundary MPEs (all-entry or no-entry), with the limit determined by initial conditions. At the boundary $\rho^* = 1$, the interior fixed point is neutrally stable and convergence is slower than geometric.*

The monotone convergence to a boundary follows from the fact that the best-response map is strictly increasing and continuous on $[0,1]$ under complementarity, which rules out cycles and chaotic dynamics. This result highlights that increasing sophistication is not universally stabilizing. Under complementarity, cognitive advancement pushes the system toward a boundary equilibrium, and the long-run outcome is path-dependent.

Remark 1 (General-case extension)**.** In the general formulation with multi-dimensional actions or continuous productive states, the scalar object J^* is replaced by a bounded linear operator on the policy-function space, obtained as the Fréchet derivative of the best-response operator at σ^*. Under Assumption 1 (C^2 payoffs, compact state space, Lipschitz kernel, and interior solutions), this operator is well defined via the implicit function theorem applied to the first-order condition of the Bellman equation, and its spectral radius $\rho(J^*)$ plays the same role as the scalar ρ^*: it governs contraction after passage to an equivalent operator norm (Stokey and Lucas 1989). Powers $(J^*)^\kappa$ denote κ-fold composition, not higher derivatives. All three propositions in this subsection carry

over under this reinterpretation, with geometric rates replaced by operator-norm rates. The scalar binary entry game furnishes all the qualitative predictions that drive the paper's main results, so the remainder of the analysis proceeds in that setting.

3.3 Evolutionary Dynamics Along the Cognitive Ladder

The distribution of cognitive states in the industry, denoted $\mu_t(\kappa) = \mathbb{P}(\kappa_{it} = \kappa \mid i \in \mathcal{N}_t)$, evolves under three forces that operate simultaneously. First, firms age deterministically along the cognitive ladder (i.e., $\kappa' = \kappa + 1$). Second, selection operates through exit, since survival probabilities depend on κ. Third, compositional replacement arises from entry, as new firms arrive with cognitive states drawn from F_0. The interaction of these forces governs how heterogeneity in strategic sophistication emerges and dissipates over time.

Lemma 3 (Cognitive Dispersion Dynamics)**.** *Under the deterministic transition $\kappa' = \kappa + 1$, the evolution of dispersion depends critically on whether new entrants are present. With ongoing entry and any initial distribution F_0, the variance $\mathrm{Var}(\mu_t)$ is non-monotone in time: it increases initially because older cohorts move up the ladder while new entrants continue to arrive at low levels, creating widening between-cohort dispersion; as more firms approach the upper bound $\bar{\kappa}$, the distribution compresses and variance declines toward zero. In contrast, under a fixed cohort (i.e., when no new entry occurs), the distribution shifts upward rigidly so that variance remains constant (equal to $\mathrm{Var}(F_0)$) until the cap begins to bind, and then declines monotonically as mass accumulates at $\bar{\kappa}$.*

The constant-variance phase is non-degenerate provided the support of F_0 lies strictly below $\bar{\kappa}$; when F_0 has mass at $\bar{\kappa}$, the cap binds immediately and variance is strictly decreasing from $t = 1$.

The lemma highlights a key implication of the cognitive ladder. Heterogeneity in strategic sophistication is not static. It is generated endogenously by the coexistence of new entrants and maturing incumbents, and it disappears once all firms become sufficiently sophisticated. This transient dispersion will play a central role in shaping both behavior and selection.

The next result characterizes how this evolving heterogeneity affects survival decisions. The key tension is that increasing cognitive sophistication has two opposing effects. It improves decision quality, but it also makes firms more forward-looking, which may reduce their willingness to remain in the market.

Assumption 3 (Positive Foresight Condition)**.** There exists $\delta > 0$ such that for all $\kappa < \bar{\kappa}$ and all x satisfying $V^{\text{true}}(x, \kappa) > \underline{V} - \bar{\varepsilon}$ (i.e., x in the region where exit is non-trivially stochastic),

$$|\hat{V}^{(\kappa)}(x) - V^{\text{MPE}}(x)| \geq \delta(\bar{\kappa} - \kappa),$$

where $\hat{V}^{(\kappa)}(x)$ is the level-κ subjective continuation value and $V^{\text{MPE}}(x)$ is the equilibrium value. This condition requires that the magnitude of the bias in continuation values is sufficiently large when firms are far from full sophistication.

The linear lower bound in $(\bar{\kappa} - \kappa)$ reflects a natural property of sequential strategic learning: the gap between a cognitively immature firm's perceived continuation value and the equilibrium value shrinks roughly in proportion to the steps remaining on the ladder. A first-time restaurateur opening in an unfamiliar neighborhood perceives the long-run value of staying through a simple projection of current demand, and that projection is far from equilibrium precisely because the restaurateur has not yet internalized the competitive responses of nearby rivals and the slow revelation of thin local demand. A restaurateur with five years of experience retains some of this optimism but materially less. A veteran operator with decades in the business holds a nearly rational view. The condition formalizes this stylized pattern by requiring that the perceptual gap be bounded away from zero in proportion to the remaining distance to full sophistication, which is needed to ensure that the intermediate-κ peak in continuation value is economically meaningful rather than vanishingly small.

Proposition 5 (Non-Monotone Survival Hazard on Cognitive State)**.** *Under Assumptions 1 and 2, the contraction condition $\rho^* < 1$, and the positive foresight condition, the conditional survival hazard $h(\kappa \mid x)$ is U-shaped in cognitive state κ for any productive state x in the stochastic exit region.*

At low levels of κ, firms make systematic strategic mistakes because they underestimate competitive crowding, which depresses realized payoffs and elevates exit. At high levels of κ, firms approach fully rational play: their subjective continuation value converges to V^{MPE}, and their exit hazard converges to the MPE-baseline rate that would obtain in a mature industry. The U-shape is therefore best read as a *trough in the middle* rather than as symmetric peaks at the ends. Intermediate firms enjoy a survival advantage because improved decision quality lifts realized payoffs

while residual optimism about continuation values keeps them above the exit margin on marginal draws of the idiosyncratic shock. Two symmetric forces produce this trough: as κ rises from zero, decision quality improves and exit falls; as κ rises toward $\bar{\kappa}$, realistic anticipation of future competition erodes the optimism premium and exit returns to the MPE baseline. Survival is consequently maximized at an intermediate level of sophistication, where both forces are partially engaged.

The mechanism depends on firms being forward-looking. If firms were myopic ($\beta = 0$), continuation values would drop out of the exit calculus entirely; exit would respond only to current-period payoffs, which are monotonically improving in κ, and the hazard would be monotonically decreasing rather than U-shaped. The intermediate peak in $V^{\text{true}}(x, \kappa)$ is thus generated by the interaction of forward-looking behavior with the evolving subjective model of rivals, not by decision quality alone. This distinguishes the mechanism from any account in which sophistication affects only current-period optimization.

The same oscillatory mechanism that generates non-monotone selection also produces non-monotone entry behavior across cognitive levels. Under strategic substitution, the level-κ entry sequence $\{p_\kappa\}$ oscillates around p^* with decreasing amplitude: odd levels over-enter relative to p^* while even levels under-enter (or vice versa, depending on p_0). The resulting pattern across cognitive levels is non-monotone, driven by these damped oscillations rather than by a clean monotone transition from over- to under-entry. High-κ firms converge to p^* as $\kappa \to \bar{\kappa}$, so the limiting behavior is neither systematic over- nor under-entry but equilibrium play.

Remark 2 (Cognitive Path Dependence)**.** The long-run equilibrium reached by the CLE depends on initial conditions through two channels. The first is standard and operates through selection in the productive state. The second is novel and operates through the evolution of cognitive states: in environments with multiple equilibria, different initial distributions of cognitive states can lead to different limiting equilibria even when the underlying primitives are identical. The sequence of cognitive states shapes the trajectory of play, which in turn determines which firms survive and which equilibrium basin the industry enters. Early differences in sophistication therefore have persistent effects that do not wash out over time.

Figure 2 illustrates the mechanism. Panel (a) plots the true continuation value $V^{\text{true}}(x, \kappa)$; Panel (b) maps this into the exit hazard $h(\kappa \mid x) = F_\varepsilon(\underline{V} - V^{\text{true}}(x, \kappa))$. Since the hazard is a

monotone decreasing transformation of V^{true}, the single-peaked value function translates directly into the U-shaped exit profile.

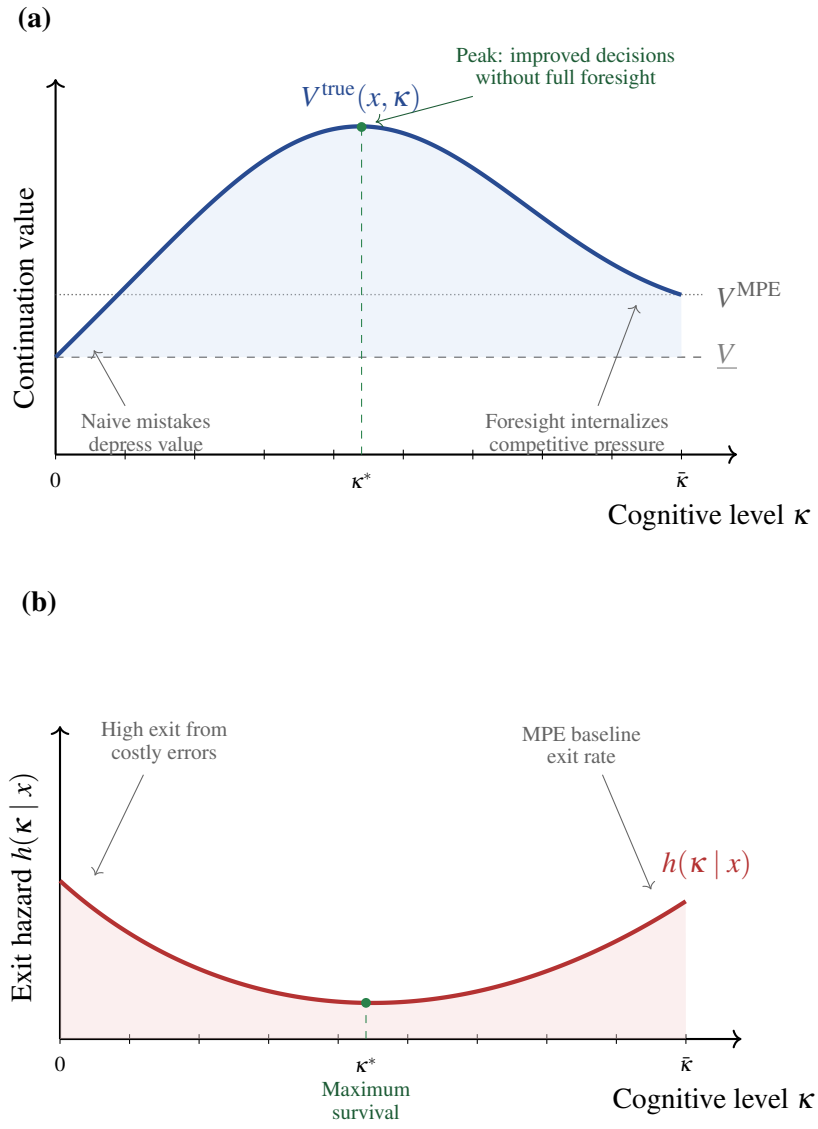
Taking the U-shape to data. The U-shaped hazard in cognitive state is directly testable once κ is proxied by observables. Natural candidates, discussed formally in Proposition 11, include firm age relative to market age, managerial tenure and prior-industry experience, repeat-entry status, and adoption of analytics or AI decision tools. Under any of these proxies, the model predicts that *exit hazard is lowest at intermediate levels of experience and elevated at both extremes* — early-stage firms exit through strategic mistakes (the left arm of the U), while mature firms exit at the MPE baseline rate once the optimism premium has been extinguished (the right arm). The prediction is non-monotone cohort survival with a trough at intermediate experience, controlling for observable productivity. Crucially, the mechanism is *discriminating*: the discussion above established that the U-shape disappears under myopia ($\beta = 0$), so observing a monotone hazard in proxied cognitive state would falsify the cognitive-ladder interpretation in favor of a learning-about-own-type explanation in the tradition of Jovanovic (1982). The U-shape is therefore not a descriptive regularity to be matched but a specific prediction that discriminates the mechanism from neighboring models with similar reduced-form survival patterns. An applied researcher with firm-age or experience data can test this directly; Proposition 11 below provides the formal framework for doing so.

3.4 The AI Acceleration Experiment

The widespread deployment of AI-powered strategic decision tools can be interpreted as accelerating movement along the cognitive ladder. In the model, this corresponds to a shift in the transition process so that firms reach the maximal level of sophistication more quickly (i.e., the maturation *horizon* T_{AI} shrinks from T_{nat} to $T_{AI} \ll T_{\text{nat}}$). Rather than gradually learning to reason strategically, firms effectively jump ahead to high levels of sophistication.

This interpretation is consistent with a growing literature showing that AI systems increasingly perform tasks requiring strategic reasoning, including anticipating rivals' behavior, adapting to feedback, and shaping competitive outcomes. Recent work conceptualizes strategy itself as a hierarchy of cognitive tasks, in which more sophisticated behavior requires deeper causal and counterfactual reasoning, and argues that advances in AI expand the range of such tasks that can

Figure 2: Non-Monotone Selection Along the Cognitive Ladder



Notes: Panel (a) shows the true continuation value $V^{\text{true}}(x, \kappa)$ as a function of cognitive level κ, for a fixed productive state x in the stochastic exit region. At low κ, naive strategic errors depress realized value below V^{MPE}. At high κ, the firm correctly anticipates long-run competitive pressure, pulling its continuation value back toward V^{MPE}. At intermediate κ^*, improved decision-making raises value while incomplete foresight prevents the firm from fully internalizing future competition, generating a single peak. The dashed line marks the scrap value $\underline{V}$; exit probability increases as V^{true} approaches $\underline{V}$. Panel (b) shows the implied exit hazard $h(\kappa \mid x) = F_\varepsilon(\underline{V} - V^{\text{true}}(x, \kappa))$. Exit is elevated at low κ because naive mistakes erode value; it returns to the MPE baseline at high κ as foresight is fully realized; and it dips to a minimum at intermediate κ^* where residual optimism about continuation values keeps firms above the exit margin. The U-shape is thus a trough in the middle, not symmetric peaks at the ends.

be automated or delegated (Chatterji et al. 2026). At the same time, evidence from multi-agent learning demonstrates that agents capable of reasoning about others' beliefs and learning processes can fundamentally alter equilibrium outcomes, sustaining cooperation or coordination even in competitive environments (Foerster et al. 2017, Yu et al. 2022). Empirical and experimental work further shows that algorithmic decision-making can reshape market behavior, including fa-

cilitating tacit coordination or acting as a strategic signal that aligns firms' actions without explicit communication (Brown and MacKay 2023, Deshpande and Jacobson 2026, Rebholz et al. 2025).

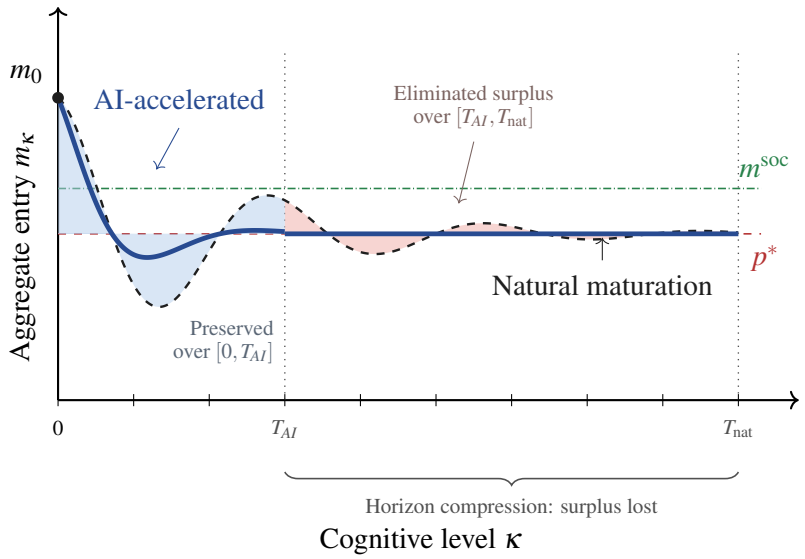
Taken together, these results suggest that AI does not merely improve optimization or forecasting, but changes the effective depth of strategic reasoning in the market. In this sense, AI adoption can be viewed as a shift in the distribution of cognitive states across firms, compressing the transition from naive to sophisticated behavior and altering the path through which industries approach equilibrium.

Scope of the horizon-compression interpretation. The framework treats AI as compressing the same cognitive ladder firms would otherwise climb naturally. This mapping is most natural for *AI-augmented* decision making: human managers use algorithmic tools to reach strategic insights faster than they would through unaided experience, but the resulting reasoning still resembles iterated best-responding of the kind the cognitive hierarchy framework describes. Two alternative interpretations lie outside the current model. First, AI might alter the *rungs* of the ladder rather than the climbing speed: algorithmic agents may reason about higher-order beliefs (Foerster et al. 2017, Yu et al. 2022) or sustain cooperation through reward-shaping mechanisms (Calvano et al. 2020) that do not fit the iterated best-response structure. Second, AI might bypass the ladder entirely, enabling direct computation of equilibrium strategies from payoff primitives — in which case the transition to p^* would be instantaneous and the transition-surplus mechanism would vanish by assumption rather than survive at a compressed horizon. My welfare results apply to the first case (AI-augmented human decision-making) and to intermediate cases where AI accelerates but does not replace iterated strategic reasoning. Settings in which AI enables genuinely different reasoning architectures — including recent empirical findings on algorithmic collusion by large language models (Fish et al. 2024) — require an extension in which the equilibrium concept itself adjusts, which I leave to future work. The horizon-compression case is of independent interest because it is empirically relevant: most organizational AI adoption to date takes the form of accelerating human decisions rather than replacing them.

Figure 3 illustrates this horizon-compression mechanism by overlaying the natural maturation path from Figure 1 against its AI-accelerated counterpart. Both paths share the same initial overentry at m_0 and converge to the same fixed point p^*, but under AI adoption the transition completes

at $T_{AI} \ll T_{\text{nat}}$. The shaded region between the natural path and p^* over the interval $[T_{AI}, T_{\text{nat}}]$ captures the transition surplus that AI acceleration eliminates: the welfare content of the slower path between the time AI-adopters have already converged and the time unaided firms would have caught up. This eliminated surplus is the object $\Phi \cdot \int_0^{T_{AI}} S(t)\,dt$ formalized in Proposition 6 below.

Figure 3: Horizon Compression under AI Acceleration



Notes: The dashed black curve is the natural maturation path from Figure 1: aggregate entry oscillates around p^* with geometrically decaying amplitude until full convergence at $T_{\text{nat}} = \bar{\kappa}$. The solid blue curve is the AI-accelerated path, which follows the same oscillatory form but with a decay rate accelerated by a factor of T_{nat}/T_{AI} (here 3), so that convergence completes at T_{AI} rather than T_{nat} and the path remains at p^* thereafter. The light blue region over $[0, T_{AI}]$ represents the transition surplus preserved under both paths. The pink region over $[T_{AI}, T_{\text{nat}}]$ represents the transition surplus eliminated by acceleration — the welfare content of the natural path's oscillation between the time AI-adopters have reached p^* and the time unaided firms would otherwise have caught up. Proposition 6 and the welfare ratio Φ formalize this eliminated-surplus quantity.

A continuous-time approximation for welfare aggregation. The model is formulated in discrete time (Section 2.1), with $t \in \{0, 1, 2, \ldots\}$ indexing periods and $T_{\text{nat}} = \bar{\kappa}$ the integer number of cognitive steps. For the welfare-aggregation and optimal-horizon analyses that follow, we adopt a continuous-time approximation in which discrete per-period surplus $S_t = W(m_t) - W(p^*)$ is replaced by a smooth envelope $S(t) = C_W e^{t \ln \rho^*}$, where the amplitude $C_W > 0$ is the first-order Taylor coefficient of W at p^* (derived explicitly in the proof of Proposition 6), and sums $\sum_t S_t$ are replaced by integrals $\int S(t)\,dt$. We extend T_{AI} from its native integer domain $\{0, 1, \ldots, \bar{\kappa}\}$ to $\mathbb{R}_+$ via this continuous envelope, so the optimal horizon T_{AI}^{opt} is a real number and the implementable horizon is its integer rounding $\lceil T_{AI}^{\text{opt}} \rceil$; Remark 3 bounds the approximation error. This approximation is innocuous for the comparative-statics exercises pursued here because (i) under strategic substi-

tution the discrete surplus alternates in sign with geometrically decaying amplitude, so pairs of consecutive periods yield a positive two-period net surplus and the continuous envelope preserves the sign and decay rate of the discrete aggregate, and (ii) the cognitive horizon $T_{\text{nat}} = \bar{\kappa}$ is large enough in settings of interest that rounding a continuous optimum T_{AI}^{opt} to the nearest integer introduces at most a constant-size distortion. The continuous-time objects $S(t)$, Φ, $\Omega(T_{AI})$, and T_{AI}^{opt} should therefore be understood as smooth approximations to intrinsically discrete quantities, not as primitive continuous-time objects.

Social discounting convention. The welfare comparisons in Propositions 6 and 9 evaluate the integrated transition surplus without social discounting. Firms themselves discount at $\beta \in (0,1)$ in their private problems, but the planner's objective is taken as the undiscounted integral $\int_0^{T_{\text{nat}}} S(t)\,dt$. This reflects the standard Ramsey position that a benevolent planner should not weight present generations above future ones within the same industry life-cycle, especially over horizons short enough that differential pure time preference is implausible as a normative criterion (Ramsey 1928). The direction of the welfare results is in any case robust to introducing a strictly positive social discount rate $\rho_s > 0$: the integrated surplus from T_{AI} to T_{nat} is still positive, so compressing the horizon still reduces welfare; only the quantitative threshold defining $\Phi > 1$ shifts from $\ln 2/(-\ln \rho^*)$ to $\ln 2/(\rho_s - \ln \rho^*)$. The sign of every comparative static in Proposition 9 is likewise preserved because ρ_s enters the welfare gradient symmetrically across T_{nat} and T_{AI}.

Remark 3 (Approximation error)**.** The approximation error can be made explicit. The pointwise identity $(\rho^*)^t = e^{t \ln \rho^*}$ holds at every integer t, so the difference between discrete and continuous cumulative surplus reduces to the standard Riemann-sum-to-integral gap for a monotone exponential integrand. A routine application of the Euler–Maclaurin bound gives $\left|\sum_{t=0}^{T_{\text{nat}}} S_t - \int_0^{T_{\text{nat}}} S(\tau)\,d\tau\right| = O\big((\rho^*)^{T_{\text{nat}}}/(1-\rho^*)\big)$, which decays geometrically in T_{nat} and vanishes as $T_{\text{nat}} \to \infty$. All comparative statics derived from the continuous-time formulation are therefore robust to the approximation at first order, with correction terms of the same geometric order as the quantities being approximated.

Proposition 6 (Acceleration Reduces Transition Surplus)**.** *A discrete reduction in the maturation horizon from T_{nat} to $T_{AI} < T_{nat}$ strictly reduces aggregate welfare relative to the natural path. The*

result holds when (a) level-0 behavior generates over-entry ($m_0 > p^*$) and (b) the MPE under-provides entry relative to the social optimum ($p^* < m^{soc}$), so that cognitive heterogeneity during the transition creates a welfare surplus. In this case, AI acceleration eliminates the surplus over $[T_{AI}, T_{nat}]$ (i.e., the transition suppression effect) while retaining only the surplus realized over $[0, T_{AI}]$.

The transition phase is not merely a path toward equilibrium but a period with independent economic value. When firms differ in sophistication, their mistakes and partial anticipation of rivals create additional competitive pressure. Once all firms behave fully strategically, this heterogeneity disappears. Accelerating cognition therefore removes part of the competitive environment rather than simply improving it.

The net welfare effect of acceleration is governed by the ratio

$$\Phi \equiv \frac{\int_{T_{AI}}^{T_{\text{nat}}} S(t)\,dt}{\int_{0}^{T_{AI}} S(t)\,dt}, \tag{4}$$

where $S(t) \equiv W^{\text{CLE}}(t) - W^{\text{MPE}}$ is the transition surplus, understood throughout as the continuous-time envelope introduced above. Because the cumulative surplus is positive throughout the transition, compressing the horizon to T_{AI} always reduces total welfare relative to the natural path. Φ measures what fraction of surplus is eliminated: when $\Phi > 1$, the surplus lost from shortening the transition exceeds the surplus retained in $[0, T_{AI}]$, indicating that the welfare cost of acceleration is large; when $\Phi < 1$, less than half of the total surplus is lost, indicating a modest welfare cost. In neither case does acceleration become welfare-improving relative to the T_{nat} baseline.

Closed-form derivation in the binary entry game. To illustrate the mechanism, consider the binary entry game with linear inverse demand $P(Q) = A - BQ$. Social welfare is

$$W(m_t) = AN_t m_t - \tfrac{B}{2}(N_t m_t)^2 - N_t m_t c_e,$$

and the socially optimal entry rate is $m^{\text{soc}} = (A - c_e)/(BN_t)$. Around the equilibrium p^*, a first-order approximation yields $S(t) \propto (m_t - p^*)$. Since the discrete-time convergence satisfies $m_t - p^* \propto (\rho^*)^t$, and using the continuous-time approximation $(\rho^*)^t = e^{t \ln \rho^*}$ with decay rate $-\ln \rho^* > 0$

(since $\rho^* \in (0,1)$), the surplus is front-loaded:

$$S(t) \propto e^{t \ln \rho^*}.$$

Substituting into (4) gives

$$\Phi = \frac{e^{T_{AI} \ln \rho^*} - e^{T_{\text{nat}} \ln \rho^*}}{1 - e^{T_{AI} \ln \rho^*}}. \tag{5}$$

Threshold condition. The welfare cost exceeds the retained surplus (i.e., $\Phi > 1$) when

$$2e^{T_{AI} \ln \rho^*} > 1 + e^{T_{\text{nat}} \ln \rho^*}.$$

In the limiting case where $(-\ln \rho^*) \cdot T_{\text{nat}}$ is large, this simplifies to $T_{AI} < \ln 2/(-\ln \rho^*)$. The key implication is that welfare losses from aggressive acceleration are most severe when strategic feedback is strong (i.e., when ρ^* is close to one, so $-\ln \rho^*$ is small, the transition surplus decays slowly, and the threshold $\ln 2/(-\ln \rho^*)$ is large). When ρ^* is close to zero (fast convergence, large $-\ln \rho^*$), the transition surplus decays rapidly, the threshold is small, and aggressive compression of the horizon would need to be extreme to generate a large welfare loss.

Note that the closed-form expression relies on a first-order approximation around p^*, which is accurate when deviations are moderate: specifically, when $|m_0 - p^*|$ is small enough that the quadratic term in the Taylor expansion of $W(m)$ is dominated by the linear term over the range $[m_0, m_0(\rho^*)^2]$ containing the first two cycles of oscillation. A sufficient condition is $|m_0 - p^*| \cdot |W''(p^*)/W'(p^*)| < 2$, which holds generically in the binary entry game with interior $p^* \in (0,1)$ for any m_0 in a neighborhood of size $O(1)$ around p^*. Under this condition, the two-period net surplus is uniformly positive and the continuous-time envelope accurately captures the discrete aggregate. The analysis also holds the number of firms fixed, treating N_t as a partial-equilibrium object. In general equilibrium, acceleration would affect entry and exit, potentially altering market structure. As shown below, these effects reinforce rather than overturn the transition suppression mechanism.

Proposition 7 (Symmetric AI Adoption Eliminates Selection on Cognition)**.** *When all firms adopt AI simultaneously, the cognitive distribution collapses immediately to $\bar{\kappa}$. This eliminates the selection mechanism operating through cognitive heterogeneity, resulting in a long-run distribution of*

productive states that is tilted toward lower-quality firms relative to the natural path.

Here, cognitive heterogeneity acts like a selection device. Low-sophistication firms exit due to mistakes, while high-sophistication firms exit due to foresight. Removing this margin preserves firms that would otherwise exit, weakening the efficiency of selection.

Definition 3 (Cognitively-Stratified Partially Oblivious Equilibrium)**.** A CSPOE describes an equilibrium in which fully sophisticated firms optimize against the evolving distribution of less sophisticated firms, while the latter continue to behave according to the cognitive ladder. The state therefore includes both fundamentals and the distribution of fringe cognitive states.

The CSPOE is characterized by a deterministic path $\{\mu_t^F\}$ generated by the operator $\mathscr{M}$, defined by

$$\mathscr{M}(\mu^F)(\kappa') = \frac{\sum_{\kappa:\min(\kappa+1,\bar{\kappa})=\kappa'} s(\kappa)\mu^F(\kappa) + \alpha F_0(\kappa')}{\sum_j s(j)\mu^F(j) + \alpha},$$

where $s(\kappa) = 1 - h(\kappa \mid x)$ and $\alpha \geq 0$ is the entrant mass. This operator maps κ to $\min(\kappa+1, \bar{\kappa})$, making $\bar{\kappa}$ the absorbing state so that fringe firms eventually reach full maturity. As time passes, the fringe distribution converges to full sophistication, and the equilibrium approaches the standard MPE.

Dominant firms condition on the current distribution μ_t^F and correctly anticipate its entire future path $\{\mathscr{M}^t(\mu_0^F)\}_{t\geq 0}$. In contrast, fringe firms treat the dominant firms' behavior as an exogenous feature of the environment and continue to apply level-κ reasoning within the fringe population.

Assumption 4 (Positive Fringe Externality)**.** The dominant firm's per-period payoff $\pi^D(x, \mu^F, a)$ is non-decreasing in the degree of fringe suboptimality. Formally, for any μ^F that is first-order stochastically dominated by the MPE distribution in the sense of *lower* aggregate fringe entry, the dominant firm's payoff satisfies $\pi^D(x, \mu^F, a) \geq \pi^{\text{MPE}}(x, a)$ for all actions a.

Assumption 4 holds most naturally when fringe immaturity causes under-entry, which softens competitive pressure and benefits the dominant firm. In the baseline over-entry specification (level-0 firms enter too aggressively), a cognitively immature fringe generates excess crowding, which would *violate* Assumption 4. The CSPOE analysis therefore applies to the late stages of the transition, when the fringe population has advanced sufficiently along the cognitive ladder that its

aggregate entry falls below the MPE level, or to alternative specifications with under-entry anchors. We flag this scope condition explicitly.

Existence. Under Assumption 4, existence of a CSPOE follows from a fixed-point argument. The dominant firm's Bellman operator is a contraction with modulus $\beta < 1$, establishing existence of an optimal policy *given* a fixed fringe path $\{\mu_t^F\}$. Because the fringe transition operator $\mathscr{M}$ depends on survival probabilities $s(\kappa) = 1 - h(\kappa|x)$, and $h(\kappa|x)$ depends on $V^{\text{true}}(x,\kappa)$ which is endogenous to the dominant firm's policy σ^D, closing the loop requires a topological fixed-point argument over the space of joint policies and transition paths. Under Assumption 1 and compactness of $\mathscr{X} \times \mathscr{M}$, a standard application of Schauder's fixed-point theorem guarantees existence of a CSPOE.

Proposition 8 (Asymmetric AI Adoption Creates Cognitive Dominance)**.** *When only a subset of firms adopts AI, these firms gain a persistent advantage over less sophisticated rivals. The resulting equilibrium (i.e., a cognitively-stratified partially oblivious equilibrium) features a dominant group that exploits the predictable behavior of a cognitively immature fringe, subject to the scope conditions of Assumption 4.*

Asymmetry in sophistication creates a systematic wedge in behavior. Sophisticated firms anticipate the evolution of the fringe and respond optimally, while less sophisticated firms continue to misperceive the environment. This generates a temporary but economically meaningful source of rents.

Remark 4 (Instability under Accelerated Complementarity)**.** Under strategic complementarity, rapid movement to high levels of sophistication can destabilize the system. When all firms become highly responsive at once, the stabilizing effect of heterogeneity disappears, and aggregate behavior converges monotonically to a boundary equilibrium rather than the interior MPE, with the path taken and the boundary reached depending on initial conditions.

Intermediate levels of sophistication dampen feedback effects under complementarity. When these intermediate stages are skipped, best responses reinforce each other more strongly, and the industry is pushed monotonically toward a boundary equilibrium that may differ from the interior MPE.

Optimal maturation horizon. The analysis highlights a fundamental trade-off between market efficiency and structural identification. Reducing the maturation horizon eliminates the beneficial heterogeneity of the transition phase, while a longer horizon prolongs a phase in which behavior is difficult to interpret and productive-state primitives are harder to identify.

Definition 4 (Identification Value)**.** Let $\nu > 0$ denote the social value per unit time of operating outside the cognitive confounding window (i.e., the value of reliable identification of productive-state primitives). The total identification value associated with maturation horizon T_{AI} is

$$I(T_{AI}) \equiv \nu \cdot \left(T_{\text{nat}} - |\mathscr{T}^{\eta}_{T_{AI}}| \right),$$

where $|\mathscr{T}^{\eta}_{T_{AI}}|$ is the length of the cognitive confounding window under the maturation horizon T_{AI}.

Proposition 9 (Optimal Maturation Horizon)**.** *Under the conditions of Propositions 2, 5, and 6, the positive foresight condition, the maintained assumption that the identification gradient $k(T_{AI}) \equiv |\partial|\mathscr{T}^{\eta}_{T_{AI}}|/\partial T_{AI}|$ is separable from ρ^* and ν in the sense that $\partial k/\partial \rho^* = \partial k/\partial \nu = 0$, and treating T_{AI} as real-valued via the continuous extension declared in Section 3.4, there exists an interior optimal maturation horizon $T^{opt}_{AI} \in (0, T_{nat})$ satisfying*

$$\left.\frac{\partial W}{\partial T_{AI}}\right|_{T^{opt}_{AI}} + \left.\frac{\partial I}{\partial T_{AI}}\right|_{T^{opt}_{AI}} = 0.$$

The optimal horizon T^{opt}_{AI} is decreasing in ν and increasing in ρ^*.

The separability assumption on k deserves comment. In the Poisson parametrization used in Corollary 1, the identification gradient depends on the shape of the score function $\partial \mu_t/\partial \tau_0$, which is driven by the cognitive advancement rate and the cap $\bar{\kappa}$ rather than by the strategic-feedback parameter ρ^*. Under the Poisson approximation, k and ρ^* enter the planner's objective additively through distinct channels — ρ^* through the welfare-gradient term $C_W e^{T_{AI} \ln \rho^*}$ and k through the identification-gradient term — which is what makes the comparative static in ρ^* tractable. If k itself depends on ρ^* (for instance, because stronger strategic feedback extends the cognitive confounding window), the sign of $\partial T^{\text{opt}}_{AI}/\partial \rho^*$ is determined by the relative magnitudes of two effects: the direct effect through surplus decay (preserves a longer horizon under strong feedback)

and an indirect effect through $\partial k/\partial \rho^*$ (ambiguous in general). The comparative static reported here isolates the direct channel; extending the result to non-separable k is a natural follow-up but requires either a full characterization of the window width as a function of ρ^* or a bound on the cross-derivative, neither of which the Poisson parametrization yields directly.

Faster maturation (smaller T_{AI}^{opt}, i.e., shorter horizon) increases informational content by shrinking the cognitive confounding window, but shortens the beneficial transition phase. When the value of identification is high, the planner prefers a shorter horizon despite the loss of heterogeneity. When competitive forces are weak (i.e., low ρ^*), the transition surplus decays rapidly, so the welfare cost of shortening the horizon is small, favoring faster adoption. Conversely, when ρ^* is close to one (strong strategic feedback, slow convergence), the surplus persists for a long time, and the planner optimally preserves a longer horizon.

The optimal horizon T_{AI}^{opt} can be interpreted as a policy target for regulating or subsidizing AI adoption. In industries where accurate productivity measurement is central to policy (i.e., high ν), a shorter horizon is desirable. In industries where such measurement plays a limited role, preserving heterogeneity may be more valuable.

3.5 Implications for Identification of Structural Models

The evolution of cognitive dispersion has direct implications for identification of structural models. Because behavior depends on the distribution of cognitive states, variation in μ_t generates variation in aggregate outcomes that is not attributable to fundamentals alone. As a result, the ability to separately identify productive and cognitive primitives depends on where the industry lies along the transition path. I develop this point as a *structural diagnostic* rather than a formal non-identification theorem: the cognitive confounding window is a conceptually tractable object that flags when standard estimators will conflate productive and cognitive primitives, and it points an applied researcher toward either (a) restricting attention to observations outside the window, or (b) adopting tools designed for games with non-equilibrium beliefs. A fully formal non-identification proof under the exact deterministic-ladder distribution is beyond the scope of this paper; the Poisson parametrization used below is a tractable approximation that captures the qualitative pattern (monotone-decaying sensitivity combined with a cap-induced pulse) without the computational burden of the exact distribution.

Corollary 1 (Non-Monotone Identification Geometry)**.** *Let $m_t(\theta)$ denote the aggregate distribution of play at time t, where $\theta = (\pi, c_e, Q_x, \tau_0)$ are structural parameters and τ_0 denotes the mean of the initial cognitive distribution under a Poisson parametrization (distinct from $\underline{\kappa}_t$ in Proposition 3, which denotes the minimum cognitive state among active firms). Under the Poisson approximation with cognitive mean $\tau_0 + t$, the sensitivity $\|\partial m_t / \partial \tau_0\|$ involves a score factor $1/(\tau_0 + t)$ that is strictly decreasing in t, reflecting that the Poisson parametrization becomes less informative over time. However, the finite cap $\bar{\kappa}$ introduces a non-monotone truncation effect: as the cognitive distribution approaches the cap, the probability mass concentrated at $\bar{\kappa}$ changes non-monotonically with τ_0, creating an identification pulse near the time of saturation. Identification of the initial cognitive distribution τ_0 thus has a non-monotone component driven by the cap-induced truncation, and vanishes as $t \to \infty$ when all mass is at $\bar{\kappa}$. AI adoption compresses this variation by shortening the maturation horizon.*

A key implication is that the source of identification difficulty is fundamentally different from standard sampling noise. The cognitive contribution to identification arises from dispersion in the population, not from idiosyncratic shocks at the firm level. In large markets, individual-level noise averages out, as in the Oblivious Equilibrium framework. In contrast, cognitive dispersion is a property of the entire population and persists even as $N \to \infty$. The resulting identification problem is therefore structural rather than statistical (i.e., it cannot be resolved by increasing sample size but only by modeling the evolution of cognition).

Proposition 10 (Cognitive Confounding Window)**.** *Define the productive-state signal ratio as*

$$\Psi_t \equiv \frac{\|J_t^x\|}{\|J_t^{\tau_0}\|}, \qquad J_t^x \equiv \frac{\partial m_t}{\partial x}, \quad J_t^{\tau_0} \equiv \frac{\partial m_t}{\partial \tau_0}.$$

For any threshold $\eta > 0$, the Cognitive Confounding Window *is the set of times*

$$\mathscr{T}^{\eta} \equiv \{t \geq 0 : \Psi_t < \eta\}.$$

When cognitive variation sufficiently dominates the variation induced by productive fundamentals, Ψ_t falls below the threshold. In this region, standard identification strategies cannot disentangle productivity from cognition. A shorter maturation horizon compresses the cognitive variation into

a shorter interval, thereby shrinking the cognitive confounding window: $\mathscr{T}^{\eta}_{AI} \subseteq \mathscr{T}^{\eta}_{nat}$.

Scope of the diagnostic. Three scope conditions deserve emphasis before drawing empirical implications. First, the window is *not point-identified* in standard datasets. In most structural IO applications, the cognitive state κ is unobserved, the cohort distribution μ_t is not directly estimable, and the Jacobians J_t^x and $J_t^{\tau_0}$ cannot be computed without a parametric model of cognitive evolution; the window's width and location are therefore not recoverable from data alone. Second, the threshold η is *a user-chosen tolerance, not a structural primitive*. Different choices of η induce different windows, and the identification-value function $v \cdot |\mathscr{T}^{\eta}|$ that enters Proposition 9 inherits this dependence. Third, the construction relies on the Poisson parametrization of Corollary 1 as a *tractability device*, not a substantive claim about the true distribution of κ in the population. The proposition should therefore be read as a *theoretical diagnostic* that identifies a failure mode of standard identification strategies in cognitively maturing industries, not as an estimator ready for off-the-shelf implementation. Operationalizing the window requires either proxy data (Proposition 11) or the non-equilibrium identification tools of Aguirregabiria and Magesan (2020).

The cognitive confounding window carries a concrete implication for empirical work. Structural estimators of industry dynamics typically identify productive-state primitives from cross-sectional and time-series variation in aggregate outcomes. The window characterizes the interval during which such variation cannot be cleanly attributed to productivity alone, because evolving cognitive dispersion is generating its own movement in aggregate play. Inside the window, estimates of productivity primitives conflate two channels; outside the window, the productivity signal is dominant and identification is standard. An analyst with access to cohort-level data can in principle avoid the window by restricting attention to observations outside $\mathscr{T}^{\eta}$. The caveat is that such data is rarely available in standard datasets used for structural IO work: the Census LBD and Compustat, for example, track firms but do not record entry cohorts in a form that cleanly partitions observations by time since market opening. The more realistic practical implication is that a researcher calibrating the width of the window for a given industry should *either* restrict estimation to mature-industry samples (where available) *or* adopt the identification tools developed by Aguirregabiria and Magesan (2020) for dynamic games with non-equilibrium beliefs, which allow payoffs and beliefs to be jointly identified under exclusion restrictions without imposing equilibrium. Their framework is more general than mine in that it accommodates arbitrary belief biases;

my contribution complements theirs by characterizing *when* the identification task is hardest (inside the cognitive confounding window, where cognitive dispersion dominates the variation their tools exploit) and by pinning down the specific bias structure that a cognitively maturing industry generates.

AI adoption compresses the path of cognitive dispersion into a shorter interval $[0, T_{AI}]$. As a result, the period during which cognitive variation dominates is both earlier and shorter, and the cognitive confounding window shrinks. This creates a second-order benefit of AI adoption by improving identification of productive fundamentals, partially offsetting the welfare loss from transition suppression. The combined effect on the planner's problem is formalized in Proposition 9: identification value $\nu \cdot (T_{\text{nat}} - |\mathscr{T}^{\eta}_{T_{AI}}|)$ enters the social objective directly, with ν parametrizing the value of operating outside the window.

Observable Proxies for Cognitive State

Cognitive state κ is not directly observable, but the identification task inside the confounding window can be sharpened by observable proxies. The following proposition establishes a sufficient condition under which a proxy restores identification of structural parameters.

Proposition 11 (Proxy-Based Identification of Structural Parameters)**.** *Let z be an observable variable satisfying:*

(R) Relevance: z is correlated with cognitive state, $\text{Cov}(z, \kappa) \neq 0$*;*

(X) Exclusion: conditional on (κ, x), z has no direct effect on the equilibrium action distribution, i.e., $\mathbb{E}[m_t \mid \kappa, x, z] = \mathbb{E}[m_t \mid \kappa, x]$.

Then inside the cognitive confounding window $\mathscr{T}^{\eta}$, structural parameters $\theta = (\pi, c_e, Q_x, \tau_0)$ are identified from the conditional moment $\mathbb{E}[m_t \mid x, z]$ provided $\text{rank}\,(\mathbb{E}[\partial m_t / \partial \theta \mid x, z])$ *equals* $\dim(\theta)$.

The proof is in Appendix 4 and proceeds by showing that conditioning on z restores separation between the productive-state Jacobian J_t^x and the cognitive Jacobian $J_t^{\tau_0}$: the first varies with x alone, the second varies with z alone (through its correlation with κ), and the joint variation recovers the full rank needed for identification. The proposition reduces the search for empirical strategies to the search for observables satisfying (R) and (X).

Several candidate proxies plausibly satisfy these conditions, ordered by increasing stringency of the exclusion restriction. *Firm age relative to market age* (calendar time since market opening) is the natural proxy in industries with a well-defined opening moment, as in the retail-entry episodes studied by Igami and Yang (2016): it is mechanically correlated with κ through the deterministic ladder and satisfies the exclusion restriction if opening dates are exogenous to firm-level productivity. *CEO or managerial experience* (prior founding of a business in a related industry) provides firm-level variation in reasoning depth; the exclusion restriction requires that prior-industry experience does not directly enhance productivity in the current industry, which is defensible when the industries share strategic structure but differ in technical production. *Analytics or algorithmic-tool adoption* shifts firms along the cognitive ladder in an observable way; here the exclusion restriction is weakest, since adoption may be correlated with omitted operational efficiency, and additional controls (e.g., firm size or prior-period productivity) are required. For the U-shaped hazard specifically, Proposition 11 combined with Proposition 5 delivers a testable prediction: *second-time entrepreneurs* (intermediate κ) should exhibit lower exit rates than both first-time entrepreneurs (low κ) and serial repeat entrepreneurs with fully-rational expectations (high κ), after controlling for observable productivity. This prediction connects the theoretical mechanism to the empirical firm-lifecycle literature in the tradition of Dunne et al. (1988) and Klepper (1996), and provides a direct falsification test of the cognitive ladder mechanism.

4 Conclusion

I have developed a theory of industry dynamics in which strategic sophistication accumulates with experience, generating a transition path from cognitively immature behavior to equilibrium play. The central insight is that this transition is not a passive adjustment process. It is itself economically meaningful, shaping competition, selection, and ultimately market structure.

Several tensions emerge from this perspective. First, there is a tension between sophistication and stability. While greater strategic ability is often associated with better decision-making, it does not necessarily stabilize markets. Under complementarity, increasing sophistication pushes the system toward a boundary equilibrium rather than dampening deviations. In such settings, the gradual progression along the cognitive ladder, by sustaining heterogeneity in behavior, may serve

as an implicit stabilizing force. The transition itself acts as a buffer against the rapid convergence to a boundary outcome that would arise if all firms behaved as fully strategic agents from the outset.

Second, there is a distinction between the private and social value of cognitive advancement. Firms that adopt more sophisticated decision tools earlier gain an advantage over less sophisticated competitors, but this advantage largely reflects redistribution rather than the creation of new value. From the perspective of the industry, the incentive to adopt such tools resembles an arms race. Each firm has an incentive to move faster, yet the collective outcome may be worse if the maturation horizon is compressed too aggressively. This highlights a misalignment between individual incentives and aggregate welfare that becomes particularly salient in the context of rapid technological change.

Third, there is a trade-off between cognitive convergence and productive selection. In the baseline dynamics, heterogeneity in strategic ability contributes to selection by differentially affecting firms' decisions and survival prospects. When all firms become sophisticated at once, this selection margin disappears. The welfare consequences depend on the correlation between sophistication and underlying productivity. If more sophisticated firms are also more productive, the natural maturation process efficiently screens out low-productivity firms through cognitive penalties; accelerating convergence eliminates this margin, preserving firms that would otherwise exit and weakening productive efficiency, consistent with Proposition 7. If sophistication and productivity are negatively correlated, natural maturation penalizes high-productivity firms, and acceleration would improve the survivor pool.

Beyond these theoretical implications, the framework generates a set of empirical predictions that can be taken to the data. The model predicts non-monotone patterns in survival over the lifecycle of a cohort, with firms at both very low and high cognitive states exiting at elevated rates for different reasons. It also implies that industries with similar fundamentals may converge to different long-run structures depending on their early cognitive composition, generating cohort-based path dependence. Finally, the analysis of AI adoption suggests that variation in the timing and speed of technological diffusion can be used to study how compressing the maturation horizon affects both competition and welfare. Although direct identification of these effects requires a structural approach, reduced-form evidence can provide informative bounds by exploiting variation in adoption across industries and over time. Proposition 11 provides the formal framework for

testing the mechanism using observable proxies for cognitive state, such as firm age relative to market age, managerial experience, and analytics-tool adoption.

Taken together, the results suggest that the evolution of industries reflects not only differences in productivity or information, but also differences in how firms reason about competition. The transition toward equilibrium behavior is itself an allocative process, and policies or technologies that compress this transition can have unintended consequences. More broadly, the framework highlights the importance of modeling strategic ability as an evolving state, providing a bridge between traditional equilibrium models of industry dynamics and behavioral approaches to strategic decision-making.

At the same time, the analysis points to several directions in which the framework can be extended. The model treats cognitive advancement as deterministic, which isolates the role of reasoning depth but abstracts from the possibility that learning itself responds to the competitive environment. Allowing the speed of cognitive advancement to depend on market conditions would introduce a feedback loop between competition and strategic ability, raising the question of whether industries converge to a stable distribution of sophistication or remain path dependent. Similarly, the model assumes a one-dimensional notion of reasoning depth, whereas in practice firms may exhibit uneven sophistication across decision margins such as entry, pricing, and investment. Introducing multidimensional cognitive states would allow for richer interactions, including complementarities and trade-offs across domains of decision-making.

A related extension concerns the scope of cognitive improvements. The framework focuses on strategic reasoning, but in practice technologies such as AI affect both strategic and non-strategic decisions. Improvements in forecasting, cost management, and operational efficiency would introduce an additional productivity channel that operates alongside the transition suppression effect identified in the paper. Understanding how these forces interact is essential for assessing the overall impact of technological change on welfare. More broadly, the framework suggests that the optimal maturation horizon is itself an economic object, shaped by the trade-off between preserving beneficial heterogeneity and improving decision quality.

References

Abbring JH, Campbell JR, Tilly J, Yang N (2018) Very simple markov-perfect industry dynamics: Theory. *Econometrica* 86(2):721–735.

Adlakha S, Johari R (2013) Mean field equilibrium in dynamic games with strategic complementarities. *Operations Research* 61(4):971–989.

Aguirregabiria V, Magesan A (2020) Identification and estimation of dynamic games when players' beliefs are not in equilibrium. *The Review of Economic Studies* 87(2):582–625.

Balbus Ł, Dziewulski P, Reffett K, Woźny Ł (2022) Markov distributional equilibrium dynamics in games with complementarities and no aggregate risk. *Theoretical Economics* 17(2):725–762.

Bartelsman EJ, Doms M (2000) Understanding productivity: Lessons from longitudinal microdata. *Journal of Economic literature* 38(3):569–594.

Besanko D, Doraszelski U, Kryukov Y, Satterthwaite M (2010) Learning-by-doing, organizational forgetting, and industry dynamics. *Econometrica* 78(2):453–508.

Brown ZY, MacKay A (2023) Competition in pricing algorithms. *American Economic Journal: Microeconomics* 15(2):109–156.

Calvano E, Calzolari G, Denicolo V, Pastorello S (2020) Artificial intelligence, algorithmic pricing, and collusion. *American Economic Review* 110(10):3267–3297.

Camerer CF, Ho TH, Chong JK (2004) A cognitive hierarchy model of games. *Quarterly Journal of Economics* 119(3):861–898.

Chatterji A, Csaszar FA, Evans J, Felin T, Hullman J, Lakhani KR, Sako M, Zenger T (2026) Can ai do strategy? a dialogue and debate. *Strategy Science* .

Decker RA, Haltiwanger J, Jarmin RS, Miranda J (2016) Declining business dynamism: What we know and the way forward. *American Economic Review* 106(5):203–207.

Deshpande S, Jacobson SH (2026) Strategic ai in cournot markets. *arXiv preprint arXiv:2601.17263* .

Doraszelski U, Satterthwaite M (2010) Computable markov-perfect industry dynamics. *RAND Journal of Economics* 41(2):215–243.

Dunne T, Roberts MJ, Samuelson L (1988) Patterns of firm entry and exit in us manufacturing industries. *The RAND journal of Economics* 495–515.

Ericson R, Pakes A (1995) Markov-perfect industry dynamics: A framework for empirical work. *Review of Economic Studies* 62(1):53–82.

Fish S, Gonczarowski YA, Shorrer RI (2024) Algorithmic collusion by large language models. *arXiv preprint arXiv:2404.00806* 7(2):5.

Foerster JN, Chen RY, Al-Shedivat M, Whiteson S, Abbeel P, Mordatch I (2017) Learning with opponent-learning awareness. *arXiv preprint arXiv:1709.04326* .

Goldfarb A, Xiao M (2011) Who thinks about the competition? managerial ability and strategic entry in US local telephone markets. *American Economic Review* 101(7):3130–3161.

Hopenhayn HA (1992) Entry, exit, and firm dynamics in long run equilibrium. *Econometrica* 60(5):1127–1150.

Ifrach B, Weintraub GY (2017) A framework for dynamic oligopoly in concentrated industries. *The Review of Economic Studies* 84(3):1106–1150.

Igami M, Yang N (2016) Unobserved heterogeneity in dynamic games: Cannibalization and preemptive entry of hamburger chains in canada. *Quantitative Economics* 7(2):483–521.

Jovanovic B (1982) Selection and the evolution of industry. *Econometrica* 50(3):649–671.

Klepper S (1996) Entry, exit, growth, and innovation over the product life cycle. *The American economic review* 562–583.

Klepper S, Graddy E (1990) The evolution of new industries and the determinants of market structure. *The RAND Journal of Economics* 27–44.

Light B, Weintraub GY (2022) Mean field equilibrium: Uniqueness, existence, and comparative statics. *Operations Research* 70(1):585–605.

Lin PH, Palfrey TR (2024) Cognitive hierarchies for games in extensive form. *Journal of Economic Theory* 220:105871.

Loewenstein G, O'Donoghue T, Rabin M (2003) Projection bias in predicting future utility. *Quarterly Journal of Economics* 118(4):1209–1248.

O'Donoghue T, Rabin M (1999) Doing it now or later. *American Economic Review* 89(1):103–124.

Pakes A (2021) A helicopter tour of some underlying issues in empirical industrial organization. *Annual Review of Economics* 13:397–421.

Ramsey FP (1928) A mathematical theory of saving. *The economic journal* 38(152):543–559.

Rebholz TR, Uphoff M, Bernges CH, Scholten F (2025) Algorithmic advice as a strategic signal on competitive markets. *arXiv preprint arXiv:2511.09454* .

Rubinstein A (1998) *Modeling bounded rationality* (MIT press).

Stokey NL, Lucas RE (1989) *Recursive Methods in Economic Dynamics* (Harvard University Press).

Weintraub GY, Benkard CL, Van Roy B (2008) Markov perfect industry dynamics with many firms. *Econometrica* 76(6):1375–1411.

Yu X, Jiang J, Zhang W, Jiang H, Lu Z (2022) Model-based opponent modeling. *Advances in Neural Information Processing Systems* 35:28208–28221.

Proof of Existence of CLE (Definition 2)

Proof. I construct the CLE by forward induction on κ. *Base case* ($\kappa = 0$). The level-0 strategy σ_0 is exogenous by assumption. *Inductive step.* Assume $\sigma^{(0)}, \ldots, \sigma^{(\kappa-1)}$ and $V^{(0)}, \ldots, V^{(\kappa-1)}$ have been constructed. At belief level κ, the opponents' action distribution $\hat{a}_{-i}^{(\kappa-1)}$ is fully determined by $\sigma^{(\kappa-1)}$ and the productive-state distribution, both of which are treated as an exogenous stochastic process by the level-κ firm under the fixed-cohort assumption. The level-κ Bellman equation therefore reduces to a dynamic program over the productive state x alone, where $\hat{a}_{-i}^{(\kappa-1)}$ is treated as a known, exogenous function of calendar time determined by the lower-level policies already constructed:

$$V^{(\kappa)}(x) = \max_{a \in \mathscr{A}} \left\{ \pi(x, \hat{a}_{-i}^{(\kappa-1)}, a) + \beta \, \mathbb{E}\left[V^{(\kappa)}(x') \mid x, a \right] \right\}.$$

Under Assumption 1, the level-κ Bellman operator is a contraction mapping on $B(\mathscr{X})$ with modulus $\beta < 1$ (Stokey and Lucas 1989), so there exists a unique fixed point $V^{(\kappa)}$ and an associated strategy $\sigma^{(\kappa)}$. By induction, this yields $\sigma^{(\kappa)}$ and $V^{(\kappa)}$ for all $\kappa = 0, \ldots, \bar{\kappa}$. The industry state process $\{s_t\}$ is then constructed by aggregating individual firm states, with consistency condition (iv) satisfied by construction. Uniqueness follows from the uniqueness of the fixed point at each induction step. □

Proof of Lemma 1

Proof. The aggregate entry rate m_t is a weighted average of level-κ entry probabilities with weights given by μ_t. By Proposition 2, under strategic substitution m_t converges to p^* with alternating sign, so the sequence $\{m_t\}$ oscillates around p^* with geometrically decaying amplitude $|m_t - p^*| \leq C(\rho^*)^t$.

Since m_0 and p^* are on opposite sides of m^{soc} by assumption, with $p^* < m^{\text{soc}} < m_0$, there exists T_0 such that $C(\rho^*)^t < m^{\text{soc}} - p^*$ for all $t \geq T_0$. For $t \geq T_0$, $|m_t - p^*| < m^{\text{soc}} - p^*$, so $m_t \in (2p^* - m^{\text{soc}}, m^{\text{soc}})$. Under strategic substitution the sequence alternates around p^*, so for infinitely many t, $m_t \in (p^*, m^{\text{soc}})$. For any such t, strict concavity of W combined with $p^* < m_t < m^{\text{soc}}$ yields $W(m_t) > W(p^*) = W^{\text{MPE}}$. Hence there exist infinitely many periods (and in particular a non-empty set of periods during the transition phase) in which $W^{\text{CLE}}(t) > W^{\text{MPE}}$. □

Proof of Proposition 2

Proof. Local contraction at the MPE fixed point. At the MPE fixed point p^*, the derivative is $J^* = -\lambda c(N-1)p^*(1-p^*)$ with $|J^*| = \rho^* < 1$. By the mean value theorem, BR is a contraction with modulus ρ^* in a neighborhood of p^*.

Level-κ recursion. The level-κ policy satisfies $\sigma^{(\kappa)} = \text{BR}(\sigma^{(\kappa-1)})$. Expanding around the fixed point σ^*:

$$\sigma^{(\kappa)} - \sigma^* = J^*(\sigma^{(\kappa-1)} - \sigma^*) + O(\|\sigma^{(\kappa-1)} - \sigma^*\|^2).$$

For p_0 in the basin of attraction of p^*, the linear term dominates for all $\kappa \geq 1$ (since for an iterated contraction around a fixed point, the quadratic remainder decays at the same geometric rate as

the linear term once the linear term initially dominates). Hence $\|\sigma^{(\kappa)} - \sigma^*\| \le C_0(\rho^*)^\kappa$ for some $C_0 < \infty$.

The sequence $\{p_\kappa\}$ converges geometrically to p^* under strategic substitution with alternating sign, so convergence is not monotone from a single side.

Aggregate play under fixed cohort. Under Assumption 2, every firm's cognitive state satisfies $\kappa_{it} = \min(\kappa_{i,0} + t, \bar{\kappa})$. For $t < \bar{\kappa}$, $\mu_t(\kappa) = F_0(\kappa - t)$ (each cohort advances by t periods). Substituting the geometric bound:

$$\|m_t - m^{\text{MPE}}\| \le C_0 \sum_{\kappa < \bar{\kappa}} F_0(\kappa - t)(\rho^*)^\kappa + \|\sigma^{(\bar{\kappa})} - \sigma^*\| \cdot \mathbb{P}(\kappa_{it} = \bar{\kappa}).$$

The first term satisfies $\sum_\kappa F_0(\kappa - t)(\rho^*)^\kappa = (\rho^*)^t \sum_j F_0(j)(\rho^*)^j \le (\rho^*)^t$. As $t \to \bar{\kappa}$, all mass accumulates at $\bar{\kappa}$ and the error floor $\|\sigma^{(\bar{\kappa})} - \sigma^*\|$ remains. Hence

$$\|m_t - m^{\text{MPE}}\| \le C_0(\rho^*)^t + \|\sigma^{(\bar{\kappa})} - \sigma^*\|.$$

Fix tolerance $\varepsilon > 0$. The saturation depth $K^*(\varepsilon)$ is the smallest κ with $\|\sigma^{(\kappa)} - \sigma^*\| \le \varepsilon$; finiteness follows from the geometric bound $\|\sigma^{(\kappa)} - \sigma^*\| \le C_0(\rho^*)^\kappa \to 0$. Choosing $\bar{\kappa} \ge K^*(\varepsilon)$ ensures the error floor is at most ε. □

Proof of Proposition 3

Proof. The OE approximation (Weintraub et al. 2008) requires that the industry state process is ergodic. In the CLE, the cognitive state distribution μ_t has a deterministic trend (each κ_{it} increases by one per period) until all firms reach $\bar{\kappa}$. This trend component does not vanish as $N \to \infty$ because it is not a finite-sample fluctuation but a deterministic function of calendar time t. The aggregate state s_t is therefore non-stationary for $t < K^*$ (the cognitive saturation depth), and conditioning on a time-invariant long-run average misspecifies the evolution of the industry. □

Proof of Proposition 4

Proof. Under complementarity, $J^* = -\lambda c(N-1)p^*(1-p^*)$, where $c < 0$ by the sign convention for complementarity (rivals' entry raises the own value of entry), so $J^* > 0$. When $\rho^* = |J^*| > 1$, the fixed point p^* is repelling: for any $p_0 \ne p^*$ in a neighborhood, the sequence $p_\kappa = \text{BR}^{\circ\kappa}(p_0)$ moves away from p^*. The best-response map $\text{BR} : [0,1] \to [0,1]$ is strictly increasing under complementarity (since $\partial \text{BR}/\partial p > 0$). A strictly increasing continuous map from $[0,1]$ to $[0,1]$ generates monotone sequences; such a map cannot produce cycles of any period or exhibit chaotic dynamics. Consequently, sequences departing from p^* converge monotonically to one of the boundary MPEs (all-entry $p = 1$ or no-entry $p = 0$), and the limiting equilibrium depends on the initial condition p_0. □

Proof of Lemma 3

Proof. Under Assumption 2, every firm's cognitive state is $\kappa_{it} = \min(\kappa_{i,0} + t, \bar{\kappa})$. Let $\sigma_0^2 = \mathrm{Var}(F_0)$ and let $\kappa_{\max}(F_0)$ denote the maximum of the support of F_0.

Phase 1 ($t \leq \bar{\kappa} - \kappa_{\max}(F_0) - 1$): variance constant. When the cap does not yet bind for any firm, $\mu_t(\kappa) = F_0(\kappa - t)$ exactly. The variance is $\mathrm{Var}(\mu_t) = \mathrm{Var}(F_0) = \sigma_0^2$, since a rigid shift preserves variance. This phase is non-degenerate provided $\kappa_{\max}(F_0) < \bar{\kappa}$; when F_0 has support reaching $\bar{\kappa}$, Phase 1 is empty and variance is strictly decreasing from $t = 1$.

Phase 2 ($t > \bar{\kappa} - \kappa_{\max}(F_0)$): variance decreases. Once the cap binds, the right tail of the shifted distribution is pooled at $\bar{\kappa}$. This pooling strictly reduces variance, and $\mathrm{Var}(\mu_t)$ converges to zero as all mass accumulates at $\bar{\kappa}$.

Ongoing entry case. The variance decomposes as:

$$\mathrm{Var}(\mu_t) = \sum_c \alpha_c \mathrm{Var}(\kappa \mid c) + \sum_c \alpha_c (\mathbb{E}[\kappa \mid c] - \mathbb{E}[\kappa])^2.$$

The between-cohort term equals zero at $t = 0$, increases as older cohorts pull ahead of new entrants, and then decreases as all cohorts converge toward $\bar{\kappa}$. Since the between-cohort term has a strictly positive interior maximum, the sum is strictly non-monotone. □

Proofs for Proposition 5 (Non-Monotone Survival Hazard)

Supporting Two-Period Example

Proof. In the two-period setting, define $V^{\text{true}}(x, \kappa) \equiv \pi(x, \bar{m}) + \beta \hat{V}^{(\kappa)}(x)$, where $\pi(x, \bar{m}) = v - c\bar{m}$ is the realized current profit and $\hat{V}^{(\kappa)}(x)$ is the level-κ subjective continuation value. This is the quantity that determines exit: the firm exits when this falls below $\underline{V} - \varepsilon_{it}$. The decomposition $\Delta V^{\text{true}}(x, \kappa) = c(\bar{m}_\kappa - \bar{m}_{\kappa+1}) - \beta[\hat{V}^{(\kappa)}(x) - \hat{V}^{(\kappa+1)}(x)]$ is then an identity.

Let m^κ denote the entry probability of a level-κ firm, and $\bar{m} = \sum_\kappa \mu_0(\kappa) m^\kappa$ the industry average. The firm's true per-period profit is $\pi(x, \bar{m}) = v - c\bar{m}$.

Level 0. $V^{\text{true}}(x, 0) < V^{\text{MPE}}(x)$, since $\bar{m} > p^*$ in a young industry, so actual profits are below equilibrium.

Level $\bar{\kappa}$. $V^{\text{true}}(x, \bar{\kappa}) = \pi(x, \bar{m}) + \beta V^{\text{MPE}}(x)$. The firm correctly anticipates future convergence.

Intermediate κ^*. $V^{\text{true}}(x, \kappa^*) = \pi(x, \bar{m}) + \beta \hat{V}^{(\kappa^*)}(x)$, where $\hat{V}^{(\kappa^*)} > V^{\text{MPE}}$ because the intermediate firm underestimates how competitive the industry will become.

Comparing: $V^{\text{true}}(x, \kappa^*) - V^{\text{true}}(x, \bar{\kappa}) = \beta[\hat{V}^{(\kappa^*)} - V^{\text{MPE}}] > 0$. So the intermediate firm has strictly higher realized value than the fully sophisticated firm. The finite differences $\Delta V^{\text{true}}(x, 0) > 0$ and $\Delta V^{\text{true}}(x, \kappa^*) < 0$ establish single-crossing, giving U-shaped $h(\kappa \mid x)$. □

Lemma 4: Diminishing Marginal Cognitive Returns

The following auxiliary lemma quantifies the rate at which marginal cognitive improvements shrink. It plays a supporting role in the intuition for Proposition 5: because decision-quality gains decay geometrically, the dominant force at high κ is the erosion of residual optimism rather than further improvements in decision quality, which is what generates the trough at an interior κ^*.

Lemma 4 (Diminishing Marginal Cognitive Returns)**.** *Under Assumptions 1 and 2 and $\rho^* < 1$, the magnitude of the decision-quality improvement $|\Delta\pi^{true}(\kappa)| \equiv |\pi^{true}(\kappa+1) - \pi^{true}(\kappa)|$ is bounded above by $C(\rho^*)^\kappa$ and converges to zero geometrically.*

Proof. Under strategic substitution, $J^* < 0$, so $p_\kappa - p^* \propto (-\rho^*)^\kappa$ alternates in sign. By Proposition 2, $|p_\kappa - p^*| \le C_0(\rho^*)^\kappa$. The improvement in realized profit from advancing one level is governed by the absolute deviation:

$$|\Delta\pi^{\text{true}}(\kappa)| = c|\bar{m}_\kappa - \bar{m}_{\kappa+1}| \le c \cdot C_0(\rho^*)^\kappa(1+\rho^*).$$

The signed quantity $\Delta\pi^{\text{true}}(\kappa)$ alternates in sign, but its magnitude is geometrically decaying at rate ρ^*. □

General Proof of Proposition 5

Proof. We use the definition $V^{\text{true}}(x,\kappa) = \pi(x,\bar{m}) + \beta\hat{V}^{(\kappa)}(x)$ established in the two-period example (which extends to the general dynamic setting by the same single-period decomposition).

Step 1: Behavior at the boundaries. At $\kappa = 0$, the firm uses the most naïve model: $\hat{V}^{(0)}(x) < V^{\text{MPE}}(x)$ when the industry over-enters (since the naive firm underestimates competitive pressure, it overestimates the value of staying). At $\kappa = \bar{\kappa}$, $\hat{V}^{(\bar{\kappa})}(x) \approx V^{\text{MPE}}(x)$ (to within tolerance ε by choice of $\bar{\kappa}$). At an intermediate level κ^*, the firm underestimates future competitive pressure, so $\hat{V}^{(\kappa^*)}(x) > V^{\text{MPE}}(x)$ by the positive foresight condition (Assumption 3).

Step 2: Single-peakedness of $V^{true}(x,\kappa)$. The cognitive state κ is integer-valued, so "single-peakedness" refers to the discrete property that there exists $\kappa^* \in \{1,\dots,\bar{\kappa}-1\}$ with $\hat{V}^{(\kappa^*)}(x) > \max\{\hat{V}^{(0)}(x), \hat{V}^{(\bar{\kappa})}(x)\}$. This follows directly from Step 1: $\hat{V}^{(0)}(x) < V^{\text{MPE}}(x)$, $\hat{V}^{(\bar{\kappa})}(x)$ differs from $V^{\text{MPE}}(x)$ by at most the saturation tolerance ε, and the positive foresight condition (Assumption 3) guarantees the existence of at least one interior κ^* with $\hat{V}^{(\kappa^*)}(x) > V^{\text{MPE}}(x) + \delta$, strictly exceeding both endpoints for $\delta > \varepsilon$. Consequently $V^{\text{true}}(x,\kappa) = \pi(x,\bar{m}) + \beta\hat{V}^{(\kappa)}(x)$ inherits this single-peaked discrete profile over $\{0,1,\dots,\bar{\kappa}\}$.

Step 3: U-shaped hazard. Since $V^{\text{true}}(x,\kappa)$ is single-peaked at κ^*, the hazard $h(\kappa \mid x) = F_\varepsilon(\underline{V} - V^{\text{true}}(x,\kappa))$ is minimized at κ^* and increases in both directions over $\{0,1,\dots,\bar{\kappa}\}$: U-shaped in κ.

Verification of the positive foresight condition in the binary entry game. In the binary entry game, the ex-ante per-period value (inclusive value) is $W(p) = \frac{1}{\lambda}\ln(1+\exp(\lambda(v-c-c(N-1)p)))$. The subjective value gap is

$$\hat{V}^{(\kappa)}(x) - V^{\text{MPE}}(x) = \frac{W(p_{\kappa-1}) - W(p^*)}{1-\beta},$$

which alternates in sign because $p_{\kappa-1}$ oscillates around p^* under strategic substitution. By the envelope theorem $W'(p) = -c(N-1)p$, the magnitude of the gap satisfies

$$\left|\hat{V}^{(\kappa)}(x) - V^{\text{MPE}}(x)\right| \approx \frac{c(N-1)p^* \cdot |p^* - p_{\kappa-1}|}{1-\beta}.$$

Because $|p^* - p_{\kappa-1}|$ is bounded *above* by $C_0(\rho^*)^{\kappa-1}$ via Proposition 2, a matching *lower* bound requires a non-degeneracy condition: we assume that the best-response iteration is non-degenerate

at the fixed point in the sense that $|p^* - p_{\kappa-1}| \geq C_0'(\rho^*)^{\kappa-1}$ for some $C_0' > 0$. This holds generically for the logit best-response map because the linear term dominates the nonlinear corrections near p^* (and fails only on a measure-zero set of payoff parameters, e.g., those producing exact cancellation between linear and higher-order terms). Under this non-degeneracy, the magnitude of the gap is at least $M \cdot (\rho^*)^{\kappa-1}$ with $M = c(N-1)p^*C_0'/(1-\beta)$. Since $(\rho^*)^{\kappa-1} \geq (\rho^*)^{\bar{\kappa}-1}$ for $\kappa \leq \bar{\kappa}$, the uniform lower bound over $\kappa \leq \bar{\kappa}$ is $M_0 \equiv M \cdot (\rho^*)^{\bar{\kappa}-1}$. Requiring $M_0 \geq \delta(\bar{\kappa} - \kappa)$ for all $\kappa \in \{0, \ldots, \bar{\kappa} - 1\}$, the binding constraint is at $\kappa = 0$ (where $\bar{\kappa} - \kappa = \bar{\kappa}$), yielding $\delta \in (0, M_0/\bar{\kappa}]$. The positive foresight condition therefore holds with a strictly positive δ, which shrinks in $\bar{\kappa}$ — an intended feature, since longer cognitive ladders require weaker per-step foresight bounds. □

Proof of Proposition 6

Proof. We derive Φ in the binary entry game with linear demand $P(Q) = A - BQ$ and sunk entry costs c_e.

Welfare function. $W(m) = AN_t m - \frac{B}{2}(N_t m)^2 - N_t m c_e$. The social optimum is $m^{\text{soc}} = (A - c_e)/(BN_t)$.

Welfare surplus. With $p^* < m^{\text{soc}}$, $\partial W/\partial m|_{p^*} > 0$. Taylor-expanding W around p^* and denoting $\Delta_t = m_t - p^*$: $S(t) \approx C_W \Delta_t$ where $C_W = N_t(A - BN_t p^* - c_e) > 0$.

By Proposition 2, $\Delta_t \approx (m_0 - p^*)(-\rho^*)^t$ under strategic substitution, with magnitude $|\Delta_t| \approx |m_0 - p^*|(\rho^*)^t$ and alternating sign. Aggregating across pairs of consecutive periods, the net two-period surplus is $C_W \Delta_t (1 - \rho^*) > 0$ when $\Delta_t > 0$, since the over-shoot amplitude exceeds the subsequent under-shoot amplitude by factor $1/\rho^* > 1$. The continuous-time approximation $(\rho^*)^t = e^{t \ln \rho^*}$ with decay rate $-\ln \rho^* > 0$ (since $\rho^* \in (0,1)$) treats $S(t)$ as a smooth front-loaded envelope of this alternating-but-net-positive discrete surplus:

$$S(t) \propto e^{t \ln \rho^*}.$$

Φ *ratio.* Substituting:

$$\Phi = \frac{\int_{T_{AI}}^{T_{\text{nat}}} C_W e^{t \ln \rho^*} \, dt}{\int_0^{T_{AI}} C_W e^{t \ln \rho^*} \, dt} = \frac{e^{T_{AI} \ln \rho^*} - e^{T_{\text{nat}} \ln \rho^*}}{1 - e^{T_{AI} \ln \rho^*}}.$$

Threshold. $\Phi > 1 \iff 2e^{T_{AI} \ln \rho^*} > 1 + e^{T_{\text{nat}} \ln \rho^*}$.

When $(-\ln \rho^*) \cdot T_{\text{nat}}$ is large, this simplifies to $T_{AI} < \ln 2/(-\ln \rho^*)$: welfare cost from acceleration is large when the compressed horizon is short relative to $\ln 2/(-\ln \rho^*)$. When ρ^* is close to one (slow convergence, small $-\ln \rho^*$), the threshold is large and many compressed horizons satisfy the condition. When ρ^* is close to zero (fast convergence, large $-\ln \rho^*$), the threshold is small, meaning only extremely aggressive compression generates a large welfare cost. □

Proof of Proposition 7

Proof. Under natural maturation, the non-monotone exit hazard of Proposition 5 generates cognitive selection: both very-low and very-high κ firms exit at elevated rates, with the exit of low-κ firms being correlated with lower x (since cognitive errors are more costly for low-productivity

firms operating near the exit margin). Under symmetric AI adoption, all firms transition to $\bar{\kappa}$ simultaneously, eliminating the period during which κ affects exit decisions. The only remaining selection is on x. The CDF of the long-run productive state distribution under AI adoption therefore first-order stochastically dominates (in the adverse direction) the distribution under natural maturation: $H_{AI}(x) \geq H_{\text{nat}}(x)$ for all x, with strict inequality in the region where low-κ firms had elevated exit. □

Proof of Proposition 8 (Profit Dominance in CSPOE)

Proof. See Section 3.4 for the CSPOE definition, the fringe transition operator $\mathscr{M}$ (with $\min(\kappa+1, \bar{\kappa})$ as the absorbing state at $\bar{\kappa}$), and the existence argument via Schauder's fixed-point theorem.

Profit dominance. Under Assumption 4, the dominant firm's best response $a^*(\mu^F)$ satisfies $\pi^D(x, \mu^F, a^*(\mu^F)) \geq \pi^D(x, \mu^F, \sigma^{\text{MPE}}(x)) \geq \pi^{\text{MPE}}(x)$, so $V^D(x, \mu^F) \geq V^{\text{MPE}}(x)$.

Non-monotone cognitive rent path. The per-period exploitation rent satisfies (to first order):

$$r_t \equiv \pi^D(x, \mu_t^F, \sigma^D(x, \mu_t^F)) - \pi^{\text{MPE}}(x) \approx C \cdot \sum_{\kappa} \mu_t^F(\kappa)(p_\kappa - p^*) + O\left(\left(\sum_{\kappa} \mu_t^F(\kappa)(p_\kappa - p^*)\right)^2\right).$$

The first-order term involves the signed net deviation $\sum_\kappa \mu_t^F(\kappa)(p_\kappa - p^*)$, which reflects cancellation effects across cognitive types: types that over-enter and types that under-enter partially offset one another. The path of r_t is non-monotone because μ_t^F first concentrates (via selective exit of extreme-κ firms) and then converges to $\bar{\kappa}$ (via maturation), with the net deviation changing sign and amplitude along the way. The total discounted rent is $\sum_{t=0}^{T_{\text{nat}}} \beta^t r_t$. □

Proof of Corollary 1

Proof. The aggregate play $m_t(\theta) = \sum_\kappa \mu_t(\kappa; \tau_0)\sigma^{(\kappa)}(\theta)$. The Jacobian component with respect to τ_0 is:

$$\frac{\partial m_t}{\partial \tau_0} = \sum_{\kappa} \frac{\partial \mu_t(\kappa; \tau_0)}{\partial \tau_0} \sigma^{(\kappa)}.$$

Under the Poisson approximation with mean $\tau_0 + t$, $\frac{\partial \mu_t(\kappa;\tau_0)}{\partial \tau_0} = \mu_t(\kappa; \tau_0) \cdot \frac{\kappa - (\tau_0 + t)}{\tau_0 + t}$ (the Fisher score). Hence:

$$\frac{\partial m_t}{\partial \tau_0} = \frac{1}{\tau_0 + t} \text{Cov}_{\mu_t}(\kappa, \sigma^{(\kappa)}).$$

The factor $1/(\tau_0 + t)$ is strictly decreasing in t. However, as the cognitive distribution approaches the cap $\bar{\kappa}$, the probability mass at $\bar{\kappa}$ responds non-monotonically to τ_0: initially it increases as more cohorts reach $\bar{\kappa}$, then stabilizes. This cap-induced truncation introduces a non-monotone component into $\partial \mu_t / \partial \tau_0$ that is not captured by the Poisson approximation. As $t \to \infty$, all mass concentrates at $\bar{\kappa}$ and $\partial \mu_t(\kappa; \tau_0)/\partial \tau_0 \to 0$ for all $\kappa < \bar{\kappa}$, so identification of τ_0 vanishes. Under AI adoption, the cap is reached sooner, compressing the identification variation into $[0, T_{AI}]$. □

Proof of Proposition 10 (Cognitive Confounding Window)

Proof. Jacobian decomposition. $J_t^x \equiv \frac{\partial m_t}{\partial x} = \sum_\kappa \mu_t(\kappa) \frac{\partial \sigma^{(\kappa)}(x)}{\partial x}$. Under Assumption 1, $\|J_t^x\| \geq c_x > 0$ uniformly in t.

Signal ratio behavior. From the identification proof, $\|J_t^{\tau_0}\|$ has a non-monotone component driven by the cap. Near the time of cap saturation, the sensitivity to τ_0 varies, and for sufficiently large η there exists a time interval where $\Psi_t < \eta$, defining a non-empty frontier $\mathscr{T}^\eta$.

AI adoption shrinks the frontier. Under AI adoption, the cap is reached at $T_{AI} < T_{\text{nat}}$. The cap-induced identification pulse is compressed into $[0, T_{AI}]$. Therefore $\mathscr{T}_{AI}^\eta \subseteq \mathscr{T}_{\text{nat}}^\eta$ for any η. □

Proof of Proposition 9 (Optimal Maturation Horizon)

Proof. Welfare gradient. From Proposition 6: $\partial W / \partial T_{AI} = S_{\text{nat}}(T_{AI}) = C_W e^{T_{AI} \ln \rho^*} > 0$. A longer horizon (larger T_{AI}) always preserves more transition surplus and thus increases market welfare.

Identification gradient. Since AI adoption compresses the variance path and shrinks the frontier: $\frac{\partial |\mathscr{T}_{T_{AI}}^\eta|}{\partial T_{AI}} > 0$, so $\frac{\partial I}{\partial T_{AI}} = -\nu \frac{\partial |\mathscr{T}_{T_{AI}}^\eta|}{\partial T_{AI}} < 0$. A longer horizon widens the cognitive confounding window, destroying identification value.

Existence of interior optimum. Define $\Omega(T_{AI}) = W(T_{AI}) + I(T_{AI})$, with derivative

$$\Omega'(T_{AI}) = C_W e^{T_{AI} \ln \rho^*} - \nu\, k(T_{AI}), \qquad k(T_{AI}) \equiv \left| \partial |\mathscr{T}_{T_{AI}}^\eta| / \partial T_{AI} \right| > 0.$$

Because $W'(T_{AI}) = C_W e^{T_{AI} \ln \rho^*}$ is strictly decreasing in T_{AI} (as $\ln \rho^* < 0$) while $\nu k(T_{AI}) > 0$, Ω' is strictly decreasing in T_{AI} whenever $k(T_{AI})$ is constant or weakly increasing over $[0, T_{\text{nat}}]$. An interior maximum requires $\Omega'(0) > 0$ and $\Omega'(T_{\text{nat}}) < 0$. At $T_{AI} = 0$: $\Omega'(0) = C_W - \nu k(0)$, which is positive iff $\nu < C_W / k(0)$. At $T_{AI} = T_{\text{nat}}$ (where $T_{\text{nat}} = \bar{\kappa}$, the number of cognitive steps, which is independent of ρ^*): $\Omega'(T_{\text{nat}}) = C_W (\rho^*)^{\bar{\kappa}} - \nu k(T_{\text{nat}})$, which is negative iff $\nu > C_W (\rho^*)^{\bar{\kappa}} / k(T_{\text{nat}})$. Both conditions hold simultaneously for $\nu \in (C_W (\rho^*)^{\bar{\kappa}} / k(T_{\text{nat}}), C_W / k(0))$, a non-empty interval whenever $(\rho^*)^{\bar{\kappa}} < k(T_{\text{nat}}) / k(0)$ (satisfied for ρ^* bounded away from one). Under these conditions, Ω' changes sign on $(0, T_{\text{nat}})$ and attains a zero at some $T_{AI}^{\text{opt}} \in (0, T_{\text{nat}})$ by the intermediate value theorem applied to the continuous-time extension of Ω declared in Section 3.4. Outside this range of ν the optimum lies at a boundary: at $T_{AI}^{\text{opt}} = 0$ when identification value is very high ($\nu > C_W / k(0)$), and at $T_{AI}^{\text{opt}} = T_{\text{nat}}$ when it is very low ($\nu < C_W (\rho^*)^{\bar{\kappa}} / k(T_{\text{nat}})$).

Comparative statics. Implicit differentiation of the first-order condition $\Omega'(T_{AI}^{\text{opt}}) = 0$, under the maintained separability assumption $\partial k / \partial \rho^* = \partial k / \partial \nu = 0$ declared in the statement of Proposition 9, yields

$$\frac{\partial T_{AI}^{\text{opt}}}{\partial \nu} = -\frac{k(T_{AI}^{\text{opt}})}{\Omega''(T_{AI}^{\text{opt}})} < 0, \qquad \frac{\partial T_{AI}^{\text{opt}}}{\partial \rho^*} = -\frac{(T_{AI}^{\text{opt}} / \rho^*) \cdot C_W e^{T_{AI}^{\text{opt}} \ln \rho^*}}{\Omega''(T_{AI}^{\text{opt}})} > 0,$$

using $\Omega''(T_{AI}^{\text{opt}}) < 0$ at an interior maximum. The first sign is intuitive: higher identification value shifts the optimum toward a shorter horizon. The second sign reflects that larger ρ^* means slower decay of the transition surplus, so more surplus is at stake over any given interval and the planner

optimally preserves a longer horizon. This comparative static is consistent with Proposition 6, which establishes that welfare losses from acceleration are most severe when ρ^* is large. □

Remark 5 (Integer implementation of the optimal horizon)**.** Because T_{AI} is integer-valued in the underlying model, T_{AI}^{opt} should be interpreted as a real-valued target horizon derived from the continuous extension declared in Section 3.4. The implemented horizon is $\lceil T_{AI}^{\text{opt}} \rceil \in \{0, 1, \ldots, \bar{\kappa}\}$, and the welfare difference between the real-valued optimum $\Omega(T_{AI}^{\text{opt}})$ and its integer rounding $\Omega(\lceil T_{AI}^{\text{opt}} \rceil)$ is bounded by the Riemann-sum-to-integral gap characterized in Remark 3: at most $O\Big((\rho^*)^{T_{AI}^{\text{opt}}}/(1-\rho^*)\Big)$, which decays geometrically in T_{AI}^{opt}. Both the existence result and the signs of the comparative statics on ν and ρ^* therefore carry over to the integer implementation, with rounding error of the same geometric order.

Proof of Proposition 1 (CLE as Restricted Optimization)

Proof. The proof proceeds in three steps.

Step 1: $\mathscr{S}^{(\kappa)}$ is well defined and closed under the best-response operator. By construction, $\mathscr{S}^{(0)}$ is the singleton containing the uniform-random policy over the action space. For $\kappa \geq 1$, define $\mathscr{S}^{(\kappa)}$ as the image of $\mathscr{S}^{(\kappa-1)}$ under the best-response operator $\mathscr{B}$, where $\mathscr{B}[\tilde{\sigma}]$ is the unique solution to the Bellman equation with all rivals fixed at $\tilde{\sigma}$. Under Assumption 1 (C^2 payoffs, compact action and state spaces, Lipschitz transition kernel), $\mathscr{B}$ is a well-defined single-valued operator by the contraction-mapping theorem applied to the Bellman equation at fixed rival strategy $\tilde{\sigma}$. Each $\mathscr{S}^{(\kappa)}$ is therefore a singleton, denoted $\{\sigma^{(\kappa)}\}$.

Step 2: $\sigma^{(\kappa)}$ solves the constrained optimization problem. Consider the constrained problem

$$\max_{\sigma \in \mathscr{S}^{(\kappa)}} \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \pi(x_t, m_t^{(\kappa-1)}) \,\middle|\, \sigma\right].$$

Since $\mathscr{S}^{(\kappa)}$ is a singleton, the constrained maximum is attained trivially at $\sigma^{(\kappa)}$. It remains to show that $\sigma^{(\kappa)}$ is also the unconstrained maximizer against the belief that rivals play $\sigma^{(\kappa-1)}$. The unconstrained Bellman equation at this belief is

$$V(x) = \max_a \{\pi(x, m^{(\kappa-1)}, a) + \beta \mathbb{E}[V(x') \mid x, a]\},$$

which has a unique solution $V^{(\kappa)}$ and induces policy $\sigma^{(\kappa)}$ by definition of $\mathscr{B}$. Hence $\sigma^{(\kappa)} = \arg\max_{\sigma \in \Sigma} \mathbb{E}[\cdots \mid \sigma]$ with the belief fixed, where Σ is the unrestricted policy space.

Step 3: Equivalence of the two formulations. The constrained formulation $\max_{\sigma \in \mathscr{S}^{(\kappa)}}$ and the unconstrained formulation with belief concentrated on $\sigma^{(\kappa-1)}$ yield the same optimizer $\sigma^{(\kappa)}$. The two formulations therefore differ only in interpretation: the former treats the strategy class as the primitive, the latter treats the belief as the primitive. Under either interpretation, $\sigma^{(\kappa)}$ is the full-optimization solution, not an approximately optimal or cost-constrained one. □

Proof of Lemma 2 (Aggregate Degeneracy of DCH)

Proof. In the DCH framework of Lin and Palfrey (2024), a level-κ player has a fixed own-level κ and maintains a posterior belief $\pi_t^{(\kappa)}$ over opponents' levels that updates within a single game as opponents' moves are observed. Own-levels do not evolve: by construction, the player's type is fixed at κ throughout the game and across any repetition of the game.

Claim 1: $\mu_t = \mu$ for all t. The population distribution of own-levels is $\mu_t(\kappa) \equiv N^{-1}\sum_{i=1}^{N} \mathbf{1}\{\kappa_i = \kappa\}$, where κ_i is firm i's own level. Since κ_i is time-invariant in DCH, $\mu_t(\kappa) = \mu_0(\kappa) \equiv \mu(\kappa)$ for all $t \geq 0$.

Claim 2 (i): Transition surplus is time-invariant. Aggregate play $m_t^{\mathrm{DCH}} = \sum_\kappa \mu_t(\kappa)\sigma_{\mathrm{DCH}}^{(\kappa)}$ depends on μ_t only, and within-game belief updating affects within-game decisions but not the own-level distribution. Under the stationary cross-game interpretation where each game is a one-shot interaction, $m_t^{\mathrm{DCH}} = m^{\mathrm{DCH}}(\mu)$ is time-invariant, and thus $W^{\mathrm{DCH}}(t) = W^{\mathrm{DCH}}(\mu)$ is time-invariant. The transition surplus $S(t) = W^{\mathrm{CLE}}(t) - W^{\mathrm{MPE}}$ in the CLE has a time-varying first term; the DCH analogue $W^{\mathrm{DCH}}(\mu) - W^{\mathrm{MPE}}$ is a single scalar.

Claim 2 (ii): U-shaped hazard is not defined on a cohort path. The exit hazard in Proposition 5 is traced by a single firm as its own cognitive state evolves from $\kappa = 0$ to $\kappa = \bar{\kappa}$. In DCH, each firm has a fixed own-level, so no such path exists. The mapping $\kappa \mapsto h(\kappa \mid x)$ is well defined cross-sectionally in DCH but does not represent the lifecycle of any individual firm, and the non-monotonic selection pattern (firms passing through the hazard trough as they age) cannot be interpreted cohort-wise.

Claim 2 (iii): Cognitive confounding window collapses. The window $\mathscr{T}^\eta$ is defined as $\{t : \Psi_t < \eta\}$ where $\Psi_t = \|J_t^x\|/\|J_t^{\tau_0}\|$ and $J_t^{\tau_0} \equiv \partial m_t/\partial \tau_0$ is the Jacobian of aggregate play with respect to the initial cognitive distribution. Since $m_t^{\mathrm{DCH}} = m^{\mathrm{DCH}}(\mu)$ is time-invariant, $J_t^{\tau_0}$ is constant in t, and Ψ_t is a constant scalar. The set $\{t : \Psi_t < \eta\}$ is therefore either empty (if $\Psi_t \geq \eta$) or the full positive real line (if $\Psi_t < \eta$). The time-variation that makes the window a useful empirical diagnostic in the CLE is absent. □

Proof of Proposition 11 (Proxy-Based Identification)

Proof. Fix $t \in \mathscr{T}^\eta$ (inside the cognitive confounding window). Without a proxy, the first-stage identification problem is to recover θ from $\mathbb{E}[m_t \mid x]$. The derivative of this moment with respect to θ decomposes as

$$\frac{\partial \mathbb{E}[m_t \mid x]}{\partial \theta} = J_t^x \cdot \frac{\partial x}{\partial \theta} + J_t^{\tau_0} \cdot \frac{\partial \tau_0}{\partial \theta}.$$

Inside $\mathscr{T}^\eta$, the signal ratio $\Psi_t = \|J_t^x\|/\|J_t^{\tau_0}\| < \eta$ is small, so variation in $\mathbb{E}[m_t \mid x]$ along θ directions that load heavily on x is obscured by the cognitive channel. Standard rank conditions fail.

Now condition on the proxy z. Under Exclusion (X), $\mathbb{E}[m_t \mid \kappa, x, z] = \mathbb{E}[m_t \mid \kappa, x]$, so z enters the conditional moment $\mathbb{E}[m_t \mid x, z]$ only through its correlation with κ:

$$\mathbb{E}[m_t \mid x, z] = \sum_\kappa \mathbb{E}[m_t \mid \kappa, x] \cdot \Pr(\kappa \mid x, z).$$

Under Relevance (R), $\Pr(\kappa \mid x,z) \neq \Pr(\kappa \mid x)$, so conditioning on z generates variation in $\mathbb{E}[m_t \mid x,z]$ that is driven entirely by variation in the cognitive weighting. Differentiating with respect to θ:

$$\frac{\partial \mathbb{E}[m_t \mid x,z]}{\partial \theta} = \underbrace{\sum_{\kappa} \frac{\partial \mathbb{E}[m_t \mid \kappa,x]}{\partial \theta} \Pr(\kappa \mid x,z)}_{\text{within-}\kappa\text{ variation, } x\text{-driven}} + \underbrace{\sum_{\kappa} \mathbb{E}[m_t \mid \kappa,x] \frac{\partial \Pr(\kappa \mid x,z)}{\partial \theta}}_{\text{across-}\kappa\text{ variation, } \tau_0\text{-driven}}.$$

The two components load on distinct structural parameters: the first on (π, c_e, Q_x) through the level-specific $\mathbb{E}[m_t \mid \kappa,x]$, the second on τ_0 through the weight distribution. Provided the rank condition stated in the proposition holds, i.e., $\operatorname{rank}(\mathbb{E}[\partial m_t/\partial \theta \mid x,z]) = \dim(\theta)$, inversion recovers θ.

The rank condition fails generically when z is weakly correlated with κ (analogous to weak-instrument concerns) or when (X) is violated and z directly shifts payoffs. The first concern is empirically testable via the first-stage regression of κ-proxies on z; the second requires institutional knowledge and overidentification tests when multiple proxies are available. □