Greedy Adaptation on an Endogenous Rugged Landscape

Huai Jie (Dante) • 29 Apr 2026

Title: Greedy Adaptation on an Endogenous Rugged Landscape: A Minimal Model of Market Inefficiency, Path Dependence, and Punctuated Equilibrium

Author: [Anonymous]

Abstract: Both market competition and biological evolution can be viewed as greedy local search algorithms on an adaptive landscape. However, markets differ fundamentally in that the landscape is *endogenous* — the payoffs to strategies deform as agents adopt them. This note constructs a tractable, stylized model that formalizes this idea. A single adaptive agent searches a rugged fitness landscape using a strictly hill-climbing rule, while the landscape itself erodes under the agent's current position due to a carrying capacity effect. We show that this minimal system exhibits three robust phenomena: (i) convergence to local optima, generating persistent inefficiency; (ii) extreme path dependence and lock-in; and (iii) punctuated equilibrium, where long periods of stasis are interrupted by sudden cascading reconfigurations. The model provides a unified micro-foundation for market inefficiency, technological lock-in, and financial crises, and it yields sharp comparative-static predictions that distinguish it from standard efficient-market theories.

—

1. Introduction

Economists have long been fascinated by the parallels between market competition and Darwinian evolution (Alchian, 1950; Nelson & Winter, 1982). Both systems reward the fit and eliminate the weak, seemingly arriving at efficient outcomes without central direction. Yet both are also subject to deep, characteristic pathologies: financial markets generate bubbles and crashes, and evolution produces fragile overspecialization and extinction. What is missing is a formal framework that explains these successes and failures with a single, precise mechanism.

We propose that the missing link is the concept of **greedy local search on a rugged fitness landscape** (Kauffman, 1993; Levinthal, 1997). In this view, an adaptive agent—whether a firm, an investor, or a biological species—

incrementally modifies its strategy and adopts only those modifications that yield an immediate fitness improvement. The landscape of payoffs is "rugged," possessing numerous local peaks. Greedy search invariably climbs one of these peaks but can never descend into a valley to reach a potentially higher summit. This alone can explain lock-in and suboptimal equilibria. In markets, however, a crucial new force arises: the landscape itself changes as the agent climbs. When many agents (or a single large agent) exploit a strategy, market impact erodes its profitability. The peak sinks.

We present a minimalist model that integrates these two mechanisms. An adaptive agent searches a rugged but *endogenous* fitness landscape via strict hill-climbing. We show that this leads to three robust, interrelated phenomena: persistent inefficiency, path dependence, and punctuated equilibrium. The model is deliberately stylized to isolate the core logic; it can be elaborated into a full agent-based or analytic framework.

2. The Model

Consider a discrete time setting, $t = 1, 2, \ldots$. There is a single representative adaptive agent that can choose a strategy from a finite set $X = \{1, 2, \ldots, M\}$. One may think of the elements of X as distinct trading rules, production techniques, or organizational forms. The agent's objective is to maximize a *realized fitness* function $\Phi(x, p_t)$ where p_t is a pollution or "crowding" variable that summarizes the history of strategy use.

Definition 1 (Rugged base landscape). The *base fitness* $f : X \to \mathbb{R}$ is a fixed mapping. We assume f is a random NK landscape: for each x, $f(x) = \frac{1}{N} \sum_{i=1}^{N} \phi_i(x_i; x_{j_1}, \ldots, x_{j_K})$, where each component contribution ϕ_i is drawn i.i.d. from a continuous distribution and depends on the state of K other components. With high K, the landscape displays many local maxima (ruggedness). The exact construction is standard (Kauffman, 1993) and need not be reproduced in full here. What matters is that f is *fixed* and *exogenous*.

Definition 2 (Erosion due to endogenous action). Let $p_t(x) \in [0, 1]$ be the *popularity* of strategy x at time t, reflecting how heavily it has been used in the recent past. This captures market impact, capacity constraints, or any diminishing-returns effect. The popularity evolves according to:

$$p_{t+1}(x) = (1 - \delta)p_t(x) + \delta \cdot \mathbf{1}_{\{x = x_t\}},$$

where $\delta \in (0, 1]$ is a memory decay parameter and x_t is the strategy chosen by the agent at time t. If the agent repeatedly uses strategy x, $p(x)$ rises toward 1; if it abandons x, $p(x)$ decays toward 0.

The *realized fitness* of strategy x at time t is then:

$$\Phi_t(x) = f(x) - \gamma \cdot p_t(x),$$

with $\gamma > 0$ measuring the strength of the endogenous erosion effect. When $\gamma = 0$, the landscape is fixed and the model is purely biological. When $\gamma > 0$, the landscape sinks under the agent's own feet—the market mechanism.

Definition 3 (Greedy local search). The agent has a *neighborhood* $\mathcal{N}(x) \subset X$, which we take to be a set of strategies that are one "mutation" away from x. In a binary string representation, $\mathcal{N}(x)$ consists of all strings differing by one bit. The agent's behavioral rule is strictly hill-climbing:

- At time t, the agent observes the current payoffs $\Phi_t(y)$ for all $y \in \mathcal{N}(x_{t-1})$ and its own payoff $\Phi_t(x_{t-1})$. - If there exists a $y \in \mathcal{N}(x_{t-1})$ with $\Phi_t(y) > \Phi_t(x_{t-1})$, the agent moves to the strategy $x_t = y$ that yields the highest such payoff (ties broken randomly). - Otherwise, the agent stays put: $x_t = x_{t-1}$. - Crucially, the agent *never moves to a strategy with strictly lower current fitness*, even if that move would open a path to a much higher peak later.

This defines a deterministic (up to ties) dynamical system on X, coupled with the popularity dynamics. The system state can be summarized by $(x_t, \mathbf{p}_t)$.

3. Analysis of Emergent Behavior

We now characterize the system's qualitative behavior. For brevity, we state results as informal propositions; formal proofs would rely on the theory of adaptive walks on NK landscapes and slow-fast dynamical systems.

Proposition 1 (Inevitable local optima). For any initial condition and any $\gamma \geq 0$, the process x_t eventually becomes trapped on a *local optimum* of the current landscape. That is, there exists T such that for all $t \geq T$, $x_t = x^*$ with $\Phi_t(x^*) \geq \Phi_t(y)$ for all $y \in \mathcal{N}(x^*)$. When $\gamma > 0$, the location of this local optimum can shift slowly as p_t evolves.

Justification. The agent only moves to higher payoffs. Since the state space is finite, a strictly increasing sequence on a bounded set must terminate. Even as Φ_t changes due to popularity erosion, the agent will only move if a neighbor becomes strictly higher; without such a neighbor, it remains locked.

Corollary 1 (Persistent inefficiency). The global optimum $x^{**} = \arg\max_x f(x)$ is typically unreachable. Since the agent never goes downhill, it cannot cross a fitness valley to reach a separate higher peak unless the landscape deforms sufficiently to eliminate the valley. Thus, the market endogenously produces and sustains inefficiency.

Proposition 2 (Path dependence and lock-in). The long-run stationary distribution of x (when it exists) is highly sensitive to initial conditions and to the sequence of early random explorations. Starting from different initial strategies, the agent ends up on different local peaks, even if the base landscape f is identical. This is a direct consequence of the greedy local search without foresight.

Proposition 3 (Punctuated equilibrium). For sufficiently slow popularity decay (small δ) and moderate γ, the system exhibits *punctuated equilibrium*: long periods of stasis around a locally optimal strategy are interrupted by sudden, rapid transitions to a new strategy.

The mechanism is as follows. Suppose the agent has settled on a local peak x^*. As long as the agent remains there, $p_t(x^*)$ increases, lowering $\Phi_t(x^*)$ relative to the fixed components of neighboring strategies. Eventually, a previously inferior neighbor y may come to satisfy $\Phi_t(y) > \Phi_t(x^*)$, triggering a jump. Once the agent moves to y, $p(x^*)$ decays while $p(y)$ rises, potentially setting off a cascade of jumps as y itself erodes. This yields sudden reconfigurations that can resemble market crashes or technological paradigm shifts, separated by long quiet periods.

Proposition 4 (Fragility and overspecialization). The system exhibits *fragility*: it becomes highly adapted to the current eroded landscape but vulnerable to a regime shift. If an exogenous shock permanently alters the base fitness landscape f (e.g., a new technology makes old peaks obsolete), the agent finds itself trapped on a low local optimum and can only adapt slowly through local steps. The greediness that made it efficient in a static world now makes it vulnerable to extinction.

4. Testable Implications and Discussion

This minimal framework yields empirical predictions that distinguish it from standard rational expectations models:

1. **Ruggedness matters:** In markets with more "complexity" (higher K in the NK landscape, i.e., many interacting factors), we should observe more persistent disequilibrium, more frequent lock-in to suboptimal technologies, and more sudden crises. 2. **Erosion intensity predicts volatility:** Higher γ (stronger market impact) leads to more frequent threshold crossings and thus to fatter tails in the distribution of strategy switches—a direct analog of volatility clustering in financial returns. 3. **Memory length shapes stability:** High δ (short memory, rapid erosion) produces more frequent, smaller adjustments; low δ (long memory) generates rare, large avalanches. This maps onto the debate about the stabilizing versus destabilizing effects of high-frequency trading. 4. **Path dependence is testable:** In lab experiments where subjects face a similar NK landscape with endogenous payoffs, different initial positions should lead to persistently different outcomes, even when the underlying fundamentals are identical.

5. Concluding Remarks

We have constructed a formal model that synthesizes the evolutionary analogy into an economic framework. The core insight is that economic systems, like biological ones, are driven by a greedy local search dynamic, but with the critical addition of an endogenous landscape—market impact that erodes the very peaks agents climb. The result is a unified explanation for the paradoxical coexistence of remarkable local efficiency and systemic fragility.

This note offers a blueprint. A full *AER*-style paper would need to fully solve the stochastic multi-agent version, characterize the invariant measure, and prove the above propositions rigorously. But the essential architecture is here: a rigorous, falsifiable, and deeply unifying theory of markets as evolution on a moving fitness landscape.

—

References (selected)

- Alchian, A. A. (1950). Uncertainty, evolution, and economic theory. *Journal of Political Economy*. - Kauffman, S. A. (1993). *The Origins of Order*. Oxford University Press. - Levinthal, D. A. (1997). Adaptation on rugged landscapes. *Management Science*. - Lo, A. W. (2004). The adaptive markets hypothesis. *Journal of Portfolio Management*. - Nelson, R. R., & Winter, S. G. (1982). *An Evolutionary Theory of Economic Change*. Harvard University Press.