Replicator Markets: When Competition Computes

Brecht Corbeel
Absolute Digital Publishers
absolutedigitalpublishers.com

August 30, 2026

Abstract

The replicator equation, the workhorse dynamic of evolutionary game theory, is not a metaphor imported into economics but a theorem about markets: the institutions a market runs — how visible rivals' performance is, how costly switching suppliers is, whether profit funds growth — fix the revision protocol its participants follow, and the protocol fixes the population dynamic that results. This paper works that mapping through in full. It unifies the demand side of market competition, customers switching suppliers, and the supply side, firms reinvesting profit into capacity, into a single two-channel replicator equation whose two intensities are separately measurable and can carry opposite signs. It proves a switching-cost damping result: a market-level selection shadow inside which quality differences smaller than the cost of switching are invisible to the dynamic. It builds an institutional taxonomy mapping five observable market dials to predicted, testable share dynamics, industry by industry, including the surprising result that deeper capital markets can make a market less evolutionary in one channel while making it more forward-looking in another. And it reads the industrial-dynamics literature's long-standing finding of weak measured selection not as a refutation of the replicator picture but as its damping term, measured.

Keywords: replicator dynamics; evolutionary game theory; imitation; revision protocols; market selection; switching costs; Walrasian selection; industrial dynamics.
JEL codes: B52, C73, D21, D83, L11, L13.
Series: The Evolutionary Economics Papers, No. 2.

1 Introduction

Two markets can look, on the surface, like instances of the same abstraction — profit-seeking firms competing for customers — and behave nothing alike. A city block of three coffee shops reshuffles market share within a season: a shop that gets noticeably worse loses walk-in customers within weeks, and a shop that gets noticeably better gains them just as fast, because anyone can taste the coffee, tell a friend, and switch tomorrow. A city's three commercial banks, or two enterprise-software vendors midway through a five-year contract, do nothing of the kind: relative quality can diverge for years while market share barely moves, because the customer relationship sits behind a switching cost, a procurement cycle, or a contract nobody is about to renegotiate early. Both markets are called "competitive" in every textbook. Only one of them is running anything that deserves to be called selection. This paper's claim is that the difference is not a matter of degree to be waved at with the word "friction" — it is a difference in which formal dynamical system the market is running, and that system is fixed, term by term, by observable features of the market's institutions.

Put as a single sentence: the replicator equation is not a metaphor imported into economics from biology, but a theorem about markets. Specific, observable market institutions — how visible rivals' performance is, how costly switching is, how tightly retained profit funds growth — implement specific revision protocols, the microscopic rule by which a customer or a firm updates its choice; revision protocols in turn determine, exactly and not approximately, the population-level dynamics those choices produce; and the replicator equation is what falls out when the protocol is imitation. The institutional design of a market therefore decides how evolutionary that market is, in a sense precise enough to attach a rate to. A companion paper in this series showed that the accounting identity Price wrote down for population genetics is, applied to firms, the same partition industrial-organization economists already use to decompose productivity growth, and treated the resulting selection intensity as a coefficient to be measured [101]. That paper never asked what fixes the coefficient's value. This one does: it derives the coefficient from a market's institutions rather than leaving it as a free parameter, unifies the demand side of competition (customers switching suppliers) and the supply side (firms reinvesting profit into capacity) in one two-channel replicator, and holds the result up against what the empirical industrial-dynamics literature has actually measured.

What is not ours

None of the machinery this claim rests on is new, and the paper states its lineage before spending a page of it. The replicator equation itself is Taylor and Jonker's [3]. Its imitation microfoundations are Schlag's [21] and Björnerstedt and Weibull's [22]. The supply-side dynamic, in which a firm's capacity grows through its own reinvested profit, is Downie's, from a 1958 book industrial economics mostly forgot [48], and Metcalfe's, who gave it its modern Fisherian form [49]. The result that imitation-driven competition in a Cournot oligopoly converges not to the Cournot–Nash quantity but to the Walrasian one is Vega-Redondo's [32], sharpened by Schaffer's finite-population account of why relative-payoff selection can favor spite over profit [33, 34] and generalized by Alós-Ferrer and Ania [35]. And the framework that organizes all of the above as instances of one construction — a revision protocol generating a mean dynamic — is Sandholm's [10]. "Markets evolve" is not a claim this paper is the first to make, and the sections that follow credit each piece of the apparatus to the person who built it, not to the paper that reuses it.

What is ours

What the paper adds, stated as exactly this and no more, is:

(a) a two-channel decomposition of market replication — demand-side switching plus supply-side reinvestment — into a single replicator equation whose two components carry separately measurable intensities, so that a firm cheap for customers but unprofitable can be shown growing through one channel while shrinking through the other (§5);

(b) the switching-cost damping proposition: a formal "selection shadow" inside which quality differences smaller than the cost of switching are invisible to the market dynamic, with the damped selection intensity given in closed form (§5);

(c) the institutional taxonomy: a systematic mapping from five observable market institutions to the revision protocols they implement, to the share dynamics those protocols predict, stated as an industry-level prediction table (§8); and

(d) a reading of the empirical audit, industry by industry, of the "weak selection" findings the productivity-growth literature has repeatedly reported: not as a refutation of the replicator picture, but as its damping term, measured (§7).

We do not claim to have discovered that markets select, that imitation produces the replicator equation, or that competitive oligopoly converges on Walrasian behavior. Those results belong to Schlag, to Björnerstedt and Weibull, and to Vega-Redondo, and are credited to them throughout. What a synthesis of this kind can add is a mapping precise enough to be wrong: given a market's institutions, a predicted dynamic, and given the dynamic, a testable rate of share convergence.

Roadmap

Section 2 traces the two histories that arrive independently at the same equation, from Maynard Smith and Price's evolutionarily stable strategy through Taylor and Jonker's dynamics on one side, and from Alchian through Nelson and Winter to Downie and Metcalfe on the other. Section 3 states the replicator equation's defining properties and its exact correspondence to the Price equation's selection term. Section 4 derives the replicator from Sandholm's revision-protocol framework and tabulates which observable protocol produces which dynamic. Section 5 states the two-channel market replicator and the switching-cost damping result, worked through a numeric example. Section 6 takes up the Walrasian-selection theorem, its spite mechanism, and its experimental scorecard. Section 7 reads the empirical industrial-dynamics literature's selection-intensity estimates against the replicator's prediction. Section 8 builds the institutional taxonomy and its industry-level prediction table. Section 9 states four places the analogy breaks — foresight, increasing returns, innovation, and the unit of selection — and what replaces the replicator there. Section 10 closes with a measurement program for the two channel intensities and this paper's place in the wider series.

2 Two histories of one equation

The replicator equation and the claim that markets select rather than optimize were discovered twice, by two literatures that did not read each other for the better part of a century. One history runs through evolutionary biology and game theory and ends with an explicit differential equation. The other runs through economics and ends, for a long time, with nothing more formal than an argument. This section tells both, in order, and stops at the point where a game-theoretic bridge was built between them without quite reaching the empirical ground either side needed.

A still earlier continuous-time adjustment process anticipated the same mechanic without the evolutionary framing that would later give it its name: Brown and von Neumann's 1950 differential- equation solution method for zero-sum games moves population mass toward strategies with a current payoff advantage, and is now read as a direct ancestor of the replicator and population-dynamic apparatus this section traces [30]. The biological history proper starts with a contest model, not a market. Maynard Smith and Price's 1973 paper on the logic of animal conflict introduced the evolutionarily stable strategy: a strategy that, once common, cannot be invaded by any rare alternative, derived without appeal to any population-level optimization principle [1]. Maynard Smith's 1982 book gave the concept its canonical book-length synthesis, formalizing the population-dynamic, proto-replicator machinery that Taylor and Jonker would shortly make fully explicit [2]. The concept was static — a condition on payoffs, not a trajectory — and it took five more years for the dynamics implied by it to be written down explicitly. Taylor and Jonker's 1978 paper is the direct mathematical ancestor of everything that follows: it derives a continuous-time population dynamic under which a strategy's frequency grows in proportion to how far its payoff exceeds the population average, and shows that ESS play is dynamically stable under it [3]. Zeeman reached a closely related dynamical classification independently and at almost the same moment, working from the qualitative theory of dynamical systems rather than from Maynard Smith's game-theoretic starting point [4], and Bomze gave a complete classification of the resulting two-dimensional phase portraits and their equivalence to the generalized Lotka–Volterra equation [15]. Schuster and Sigmund's 1983 paper gave the resulting family of equations the name it has carried since, observing that the identical structure recurs across population genetics, ecology, prebiotic chemistry, and sociobiology [5]. By the early 1980s, in other words, three independent routes — game-theoretic, dynamical-systems, and biochemical-evolutionary — had converged on the same equation, which is one reason its later economic reappearance is closer to a rediscovery than an import.

The following two decades were spent proving what that equation implies rather than finding new forms of it. The connection between the dynamic's rest points and the game's Nash equilibria, and between evolutionary stability and asymptotic stability, was assembled piece by piece into what the field now treats as settled background — synthesized comprehensively in Hofbauer and Sigmund's graduate textbook, restated in their own later survey, and restated again in Cressman and Tao's survey after that [6, 7, 13]. Section 3 states that package formally; here it matters only as the second landmark on the biological side's timeline. A third development complicated the picture rather than completing it: once mutation or experimentation is added to the deterministic dynamic, the question of which equilibrium survives in the long run becomes a question about which equilibrium's basin of attraction is hardest for noise to escape. Foster and Young gave the concept of stochastic stability that answers this [17], and Kandori, Mailath, and Rob, together with Young's independent and near-simultaneous paper on the evolution of conventions, showed that persistent low-probability experimentation selects the risk-dominant equilibrium in simple coordination games rather than whichever equilibrium the deterministic dynamic happens to start near [18, 19]. Samuelson's book-length treatment surveys how this same machinery resolves the equilibrium-selection problem across noncooperative games more broadly — drift, noise, and backward and forward induction included — rather than only in the coordination-game setting Kandori-Mailath-Rob and Young analyze [9]. This machinery — deterministic replicator flow plus a stochastic-stability refinement riding on top of it — is what Section 6 redeploys to state the paper's Walrasian-selection result.

That is one history, and it produced an equation before it had any use for one. The economics side ran on almost the opposite schedule: it had an argument decades before it had an equation, and the argument arrived already contested. Veblen's 1898 essay asking why economics was not an evolutionary science set the terms half a century early, indicting the discipline's static habits of mind and demanding an account of economic life as a cumulative, non-equilibrium process [41], but the demand sat unanswered until Alchian's 1950 paper supplied the first mechanism with teeth: firms need not calculate or maximize anything for a competitive market to behave as though they had, because it is realized profitability after the fact, not foresight beforehand, that determines which firms survive to be imitated and which disappear [42]. Enke, writing the following year on the Chamberlin–Robinson distinction, reached for the same idea from the industrial-organization side, invoking what he called viability analysis and the forces of natural selection in the economic environment to argue that survivors behave as maximizers whether or not any of them are [43]. Penrose's 1952 rejoinder is the discipline's founding objection, not a footnote to it: biological analogies applied to firms strip out exactly what makes a firm a firm rather than an organism, namely conscious managerial choice under conditions the manager can partly control, and a selection story with no intentionality in it is not a theory of the firm [44]. Friedman's 1953 methodological essay closed the debate's first round by making the as-if defense general and explicit — a theory should be judged by its predictions, and natural selection among market participants makes the profit-maximization assumption predictively adequate whatever its psychological realism [45]. Nothing in this exchange was yet a differential equation.

Winter supplied the formal turn the debate had been missing. His 1964 paper on economic "natural selection" and the theory of the firm built, independently of and close in time to the Alchian argument, an explicit model of selection acting on firms understood as bearers of decision rules rather than as optimizing agents in the conventional sense [46]. Nelson and Winter's 1982 book turned that formal starting point into a research program: organizational routines are the heritable unit, search and innovation the source of variation, and selection on realized market performance the mechanism that reweights the population of routines toward the fitter ones [47]. It is against this program's mathematical content that the present paper's own contribution is properly measured, not against the looser Alchian or Friedman arguments that preceded it. Dosi and Nelson's own survey restates the shared variation–selection–retention methodology this program established as evolutionary economics' field-defining statement, and their later synthesis two decades on updates it with the intervening evidence on technological change and industry dynamics [52, 53].

One contemporaneous strand of this history is easy to miss because it was never widely read. Downie's 1958 book *The Competitive Process* is not digitized anywhere openly reachable today, and this paper's own knowledge of it runs through a contemporaneous review rather than the primary text [48]; that fact is itself part of the point this section is making about two literatures failing to find each other. What the review confirms is that Downie modeled competition as operating directly on the dispersion of efficiency between firms, with progress — not equilibrium — as the theory's explicit object. The mechanism, on this reading, is a transfer of resources and market position away from less efficient producers and toward more efficient ones, driven by the fact that a more efficient firm's higher profit funds the very expansion that lets it take share from its rivals. That is a supply-side channel distinct from anything in the Alchian–Winter imitation story, and it is the direct historical ancestor of the reinvestment dynamic this paper states formally as part of Proposition 2 in Section 5.

Whether or not the citation trail runs continuously from Downie forward, the mechanism resurfaces, restated in explicitly Fisherian terms, in Metcalfe's work three decades later. Metcalfe's 1994 paper applies Fisher's fundamental theorem of natural selection directly to a population of competing firms: under replicator share dynamics, the rate of change in an industry's mean productivity is driven by the variance in relative fitness — market-share growth — across its constituent firms, with a chapter of his 1998 book bearing the title "Fisher's Principle and the Process of Competition" to make the identification unmistakable [49, 50]. This is the economics side's own arrival, independent of the biological literature's textbook synthesis, at a population-genetics fundamental theorem applied literally rather than metaphorically to an economy — and it is the direct precedent for the two-channel formalization this paper builds in Section 5.

A formalization wave followed, arguing that the analogy Penrose had rejected in 1952 deserved to be read as an ontology rather than a metaphor. Andersen's 2004 paper proposes what he calls "evometrics": applying the Price equation directly to populations of firms and industries to partition economic change into a selection component and an innovation component [58]. Knudsen's parallel program argues that economic selection theory should be built to mirror the causal structure of neo-Darwinian theory — variation, replication or heredity, and selection — using Alchian's and Nelson and Winter's accounts as illustrations of how organizational and market-level selection interact [59]. Hodgson and Knudsen's book-length case for "Darwin's conjecture" pushes the claim furthest: stripped of their genetic particulars, Darwinian principles are argued to be the correct general theory of any evolutionary process, economic change included [55], and the multi-author defense of "generalized Darwinism" answers the standard objection that this is reductionism by insisting variation-selection-retention is a domain-general causal schema, not a biological one on loan [56]. Safarzynska and van den Bergh's survey maps the resulting field's building blocks — diversity, selection, innovation, diffusion, path dependency, multilevel selection — across its three main formal traditions of evolutionary game theory, evolutionary computation, and agent-based modeling [57]. Witt's later methodological stock-taking distinguishes the routines-based Nelson–Winter tradition from generalized Darwinism as competing internal readings of what "evolutionary" should mean in this literature [54], and Hodgson's book-length history of the field surveys its post-1980 fragmentation across disciplines, arguing generalized Darwinism could supply the unifying core theory it still lacks [60].

The bridge that finally carried formal evolutionary game theory into economics as a general modeling tool, rather than as an argument by analogy, was built by Friedman and by Mailath, and it is worth stating plainly where it stopped. Friedman's 1991 paper introduces evolutionary games as models of repeated, anonymous strategic interaction in which more-fit actions displace less-fit ones over time, giving economists the fitness functions and compatible-dynamics apparatus needed to treat market competition itself as a game-dynamic process — the framing this paper's title takes at face value [39]. Weibull's contemporaneous textbook did the same work in a systematic, teachable form, translating the biological replicator formalism into the standard apparatus of strategic economic behavior [8]. Mailath's 1998 survey asks the pointed question directly: do people play Nash equilibrium, and can evolutionary dynamics explain which equilibrium gets played when several are available [40]. Both papers are rigorous, and both are theoretical. Neither one, nor the literature they anchored, closed the loop with the empirical industrial-dynamics program that Section 7 draws on — the firm-growth, productivity-dispersion, and selection-intensity measurements that by the 1990s already had decades of their own accumulated evidence, built by economists who were not citing Maynard Smith, Taylor and Jonker, or Metcalfe at all. Two histories, one equation, and a gap between the equation's proof and its measurement that the remaining sections of this paper are written to close.

3 The replicator equation and what it commits you to

Everything that follows is built on one equation. Let $x_i(t) \geq 0$ denote the share of strategy (or firm) i in a population with $\sum_i x_i = 1$, let π_i denote its payoff or profit rate, and let $\bar{\pi}$ denote the population's mean payoff, $\bar{\pi} \equiv \mathbb{E}_s[\pi_i] = \sum_i x_i \pi_i$. The replicator equation states that each share grows at a rate equal to its own payoff advantage over that mean:

$$\dot{x}_i = x_i\,(\pi_i - \bar{\pi}). \tag{1}$$

Nothing about Equation (1) refers to what any individual is trying to do. It is a statement about a population, written down as a claim about frequencies, and its usefulness for the rest of this paper rests on three properties that hold regardless of what π_i is taken to mean substantively — consumer surplus, firm profit, or anything else share-weighted and comparable across i.

The first property is that the simplex is invariant: if $x_i(0) \geq 0$ for all i and $\sum_i x_i(0) = 1$, then $x_i(t) \geq 0$ and $\sum_i x_i(t) = 1$ for every t thereafter [3, 6]. This is not an added assumption; it is forced by the equation's own structure. Summing Equation (1) over i gives $\sum_i \dot{x}_i = \sum_i x_i \pi_i - \bar{\pi} \sum_i x_i = \bar{\pi} - \bar{\pi} = 0$, so the total share never drifts off one, and every face of the simplex on which some $x_i = 0$ is itself invariant, since $\dot{x}_i = 0$ there — a strategy or firm with zero share cannot spontaneously reappear under this dynamic. A market-share vector stays a market-share vector for all time; this is what licenses reading x_i or s_i as a share throughout the rest of the paper rather than as an unbounded population count that merely happens to sum to one at $t = 0$.

The second property is a family of invariances under transformations of the payoffs that leave the dynamics themselves, or at least their orbits, unchanged. Because only the difference $\pi_i - \bar{\pi}$ enters Equation (1), adding any common term $c(x)$ — a function of the population state but not of i — to every payoff leaves $\dot{x}_i$ exactly as it was: a uniform shift in profitability that raises all firms' margins equally has no effect on how shares move, only a shift that changes relative standing does. More generally, replicator-type dynamics, imitation dynamics, and several other adjustment processes can be organized as different projections, or quotients, of a single underlying family, so that state-dependent rescalings of time and payoff that preserve relative fitness comparisons preserve the shape of the resulting trajectories even when they change the speed at which those trajectories are traversed [14]. Both facts matter substantively later: the two-channel replicator of Section 5 adds two structurally different payoff measures, consumer surplus and profit, into one dynamic precisely because the equation does not care about the payoffs' absolute scale or common level, only about their share-weighted spread.

The third property is the reason the equation is taken seriously as a model of strategic behavior rather than merely as a bookkeeping device for reweighting frequencies, and it is stated here as a theorem rather than derived, since its proof is the joint work of several decades of the literature reviewed in Section 2.

Theorem 3.1 (Folk theorem for the replicator dynamic)**.** *Under Equation* (1)*: every interior rest point is a Nash equilibrium of the underlying game; every Lyapunov-stable rest point is a Nash equilibrium; every strict Nash equilibrium is asymptotically stable; and every evolutionarily stable strategy is asymptotically stable.*

This is the folk theorem of evolutionary game theory, assembled and proved across the literature and stated here in the form given by Hofbauer and Sigmund's synthesis and Cressman and Tao's later survey [6, 13], building on the evolutionarily stable strategy concept of Maynard Smith and Price [1]. Read carefully, Theorem 3.1 is a set of one-way implications, not an equivalence, and the direction matters: it tells you that wherever the dynamic settles, or wherever it is pinned down by strictness, game-theoretic equilibrium is waiting, but it does not tell you that the dynamic settles anywhere in particular, nor that every Nash equilibrium is a plausible resting place. A market governed by Equation (1) is not guaranteed to converge; it is guaranteed that if and where it does converge to an interior point, or is held at a point by strict best-reply incentives, that point is an equilibrium of the game the payoffs define. That asymmetry — equilibrium as a possible resting place of a process, not as an assumption about behavior — is the whole methodological content of treating competition as a replicator system rather than as an optimization problem solved once and for all. The formalism earning that status is not a biology-only tool pressed into economic service: Nowak's wide-ranging synthesis applies the same replicator-equation machinery across games, language, cancer, and viral dynamics [11], and Nowak and Sigmund's review makes the general point explicit, arguing that replicator and adaptive dynamics, not optimization algorithms, are the appropriate calculus wherever selection is frequency-dependent [12].

The equation earns a second, independent justification from an entirely different direction: it is, to first order, an accounting identity rather than a behavioral model at all. Suppose time is discretized in steps of length Δt and each strategy's relative fitness over a step is defined as $w_i = 1 + \Delta t\,(\pi_i - \bar{\pi})$, so that a strategy with above-average payoff grows and one with below-average payoff shrinks in exact proportion to its payoff gap. Then, as Paper No. 1 in this series establishes, the discrete-time replicator dynamic this generates is not an independent postulate about market behavior; it is exactly the selection term of the Price equation with the transmission term set to zero. For any share-weighted characteristic z_i carried by strategy i — efficiency, quality, whatever the rest of the paper needs it to be —

$$\Delta \mathbb{E}_s[z] = \mathrm{Cov}_s(\pi_i, z_i)\,\Delta t + O(\Delta t^2) \tag{2}$$

[101]. The change in the population's mean z is, to leading order in the length of the revision interval, exactly the share-weighted covariance between payoff and z — Fisher's fundamental theorem in the form Price gave it, with no assumption anywhere that any individual chose z_i or is trying to raise it. Equation (2) is why this paper's later empirical audit in Section 7 can treat published productivity-decomposition covariance terms as direct measurements of the same object Equation (1) describes, rather than as evidence bearing on it only by loose analogy: the replicator dynamic and Paper No. 1's selection accounting are, at this order, the identical statement about the identical data.

None of this holds exactly once the population in question is finite, and every application in this paper is, eventually, a finite population of firms. Equation (1) is a continuum limit, and the limit is not always benign. Traulsen, Claussen, and Hauert show that the microscopic stochastic dynamics of a finite population can differ qualitatively from — in some regimes even invert the direction implied by — the deterministic replicator limit that is supposed to approximate it as population size grows [16]. Nowak, Sasaki, Taylor, and Fudenberg's finite-population analysis of cooperation gives the best-known instance of this kind of correction, the "one-third rule" governing when a rare strategy can invade under weak selection in a population of a specific finite size — a condition with no counterpart in the infinite-population dynamic at all [98]. And the discrepancy is not confined to knife-edge invasion conditions: Schaffer's finite-population extension of the evolutionarily stable strategy concept shows that the finite-N analogue of an ESS is generally *not* the Nash equilibrium of the same game, because relative-payoff comparisons in a finite population create incentives — Section 6 calls this spite — that have no effect once the population is large enough for any one player's action to leave the mean payoff unmoved [33]. Every market this paper analyzes after this section is finite. Equation (1) is still the right organizing object for it, for the reasons this section has given, but the folk theorem's Nash-equilibrium guarantees, and the finite-population corrections that qualify them, both travel together into everything that follows.

4 Revision protocols: which markets are replicator markets

Section 3 treated the replicator equation as a dynamical object with a settled pedigree and a settled folk theorem. It said nothing about when a market should be expected to obey it. That question has a precise answer, and it is not a question about markets specifically: it is a question about what a *revising agent* can see and do when a chance to change strategy arrives. **(author?)** [10] organizes the whole family of population dynamics — replicator, best-response, logit, pairwise-comparison — as instances of a single object, the *revision protocol*: a rule $\rho_{ij}(\pi, x)$ specifying the rate (or, in discrete time, the probability) at which an agent currently playing strategy i switches to strategy j, as a function of the payoff vector π and the current population state x. Different protocols encode different behavioral assumptions — imitate, best-respond, compare, reinforce — and **(author?)** [14] show the resulting dynamics nest inside one quotient/projection framework rather than constituting unrelated formalisms. The replicator equation is one point in this space, not the whole space. Locating a market in it is therefore an empirical and institutional question: what do participants observe, and what do they do with it?

Aggregating a revision protocol over a large population gives its *mean dynamic*. If a mass x_j of agents currently plays j and each such agent, on receiving a revision opportunity, switches to i at rate $\rho_{ji}(\pi, x)$, then the population share x_i gains mass $x_j\rho_{ji}$ from every source strategy j and loses mass to every destination strategy j at rate ρ_{ij} from its own x_i-mass of current adherents. Summing inflows and subtracting outflows gives Sandholm's general mean-dynamic equation:

$$\dot{x}_i \;=\; \sum_j x_j\,\rho_{ji}(\pi, x) \;-\; x_i \sum_j \rho_{ij}(\pi, x).$$

This equation is bookkeeping, not behavior: it holds for any revision protocol whatsoever, including ones with no economic content. All the economic content is in the choice of ρ. The question "which markets are replicator markets" is the question of which ρ a given institution implements.

Proposition 1 (Imitation microfoundation; 21)**.** *Under pairwise proportional imitation: a revising agent samples one other agent uniformly and switches to their strategy with probability proportional to the positive part of the payoff difference,* $\rho_{ij} = \kappa\,[\pi_j - \pi_i]_+$*. The resulting mean dynamic is exactly*

$$\dot{x}_i \;=\; \kappa\, x_i\,(\pi_i - \bar{\pi}).$$

The derivation sketch is short; Appendix A gives it with every step. Sampling "one other agent uniformly" means the probability that the sampled partner plays j is x_j, so the condi-

Revision protocol	Resulting dynamic	Informational requirement
Pairwise proportional imitation	Replicator	Observe one rival's realized payoff
Imitation of success (Björnerstedt–Weibull)	Replicator (up to speed)	Observe payoffs
Best response (Gilboa–Matsui)	BR dynamic	Know the whole payoff function
Logit / perturbed BR (Blume)	Logit dynamic	Know payoffs, noisy choice
Pairwise comparison (Smith)	Smith dynamic	Compare two candidate payoffs
Reinforcement (Börgers–Sarin, Erev–Roth)	Replicator-like in expectation	Own payoff only

Table 1: Revision protocols and the population dynamics they generate.

tional rule quoted in the proposition combines with the encounter probability to give the full state-dependent protocol $\rho_{ij}(\pi, x) = x_j\, \kappa\, [\pi_j - \pi_i]_+$, and symmetrically $\rho_{ji}(\pi, x) = x_i\, \kappa\, [\pi_i - \pi_j]_+$. Substituting both into Sandholm's general equation, factoring the common x_i out of both the inflow and outflow sums, and applying the elementary identity $[a]_+ - [-a]_+ = a$ to the bracket difference $[\pi_i - \pi_j]_+ - [\pi_j - \pi_i]_+$ collapses the sum over j to $\sum_j x_j(\pi_i - \pi_j) = \pi_i - \bar{\pi}$, leaving $\dot{x}_i = \kappa\, x_i(\pi_i - \bar{\pi})$ — the replicator equation, up to the constant κ that sets the clock speed of revision. Schlag's own contribution is prior to this population-level step: he shows the proportional rule is not an arbitrary behavioral postulate but the rule a boundedly rational agent, drawing sequential bandit-style inferences from a single sampled observation, should actually use if it wants to avoid ever choosing the worse of two options in expectation [21]. The replicator equation therefore is not merely *consistent with* individually sensible behavior; on Schlag's argument it is what individually sensible behavior aggregates to. Helbing's earlier stochastic, master-equation derivation of behavior change through pairwise contact between individuals reaches a structurally similar imitation dynamic from a physics-style Boltzmann-interaction model, independently of Schlag's bandit-inference argument [99].

Two features of Proposition 1 carry directly to the market setting. First, the informational requirement is minimal: a revising agent needs to observe one rival's *realized* payoff, nothing about the shape of the payoff function, nothing about other rivals, and nothing about why the rival is doing well. Second, the protocol says nothing distinctively economic — no prices, no profit-maximization, no market clearing enters the derivation. The replicator equation is what falls out of *any* population in which revision means "occasionally look at how someone else is doing and copy them with a probability that rises in how much better they are doing." A market is a replicator market to the extent that its institutions implement, even loosely, that rule: point-of-sale comparison shopping, franchise-performance benchmarking, review platforms and word-of-mouth, retained customers watching a competitor's line move faster, are all instances of pairwise proportional imitation running on economic strategies (which supplier to patronize, which technique to adopt) rather than on animal contests.

But proportional imitation is only one cell in Sandholm's typology, and the other cells matter because they are the ones a given market's institutions might implement *instead*. Table 1 lays out the correspondence the rest of the paper leans on: for each revision protocol, the resulting population dynamic, and what a revising agent must be able to observe for that protocol to be behaviorally available at all.

The first two rows both deliver the replicator equation, or something asymptotically equivalent to it, from imitation rules that differ in how the comparison is made. **(author?)** [22] examine imitation of the population's currently most successful strategy rather than a single random sample; this imitation-of-success protocol reproduces the replicator dynamic's direction of motion everywhere but can differ from it in speed, since the rate at which the best strategy is identified and copied need not equal the rate at which a uniformly sampled comparison happens to favor it. The remaining rows show what is lost as the informational assumption is strengthened rather than relaxed. Best response, in the tradition running from **(author?)** [27]'s fictitious play through **(author?)** [26]'s socially stable strategies, requires a revising agent to know the entire payoff function well enough to compute an optimal reply to the current population state — a demand no comparison-shopping consumer and few competing firms can meet, and the dynamic it generates (a differential inclusion jumping toward the best-reply correspondence) looks nothing like smooth replicator flow. **(author?)** [25]'s logit protocol softens best response with choice noise but keeps the requirement that the whole payoff structure, not just one rival's realized outcome, be known. **(author?)** [29]'s pairwise-comparison protocol, developed for congestion routing rather than markets, sits between the imitation and best-response families: it needs a revising agent to evaluate two candidate payoffs directly rather than merely observe one, which is a heavier requirement than proportional imitation even though the resulting Smith dynamic shares the replicator's Nash stationary points. The ordering across the table is an ordering of what a market's institutions would have to supply for each protocol to be live, and it is this ordering that the taxonomy of Section 8 turns into a testable mapping from observable market structure to predicted dynamics.

The table's last row is a second, independent route into the replicator family, and it removes even the modest requirement of observing a rival at all. **(author?)** [23] show that a Cross-type stimulus-response reinforcement learner — one who raises the probability of playing a strategy after it pays off well and lowers it after it pays off badly, using only its *own* payoff history and no information whatsoever about rivals — converges, in the continuous-time limit, to the same replicator dynamic that pairwise imitation produces at the population level. **(author?)** [24] show a one-parameter version of this rule tracks observed human play of mixed-strategy experimental games better than the Nash prediction it is meant to approximate, evidence that reinforcement of this kind is not a modeling convenience but a realistic description of how boundedly rational agents actually adjust. **(author?)** [31] ties the two routes of this section together formally: reinforcement learning and stochastic fictitious play are both perturbations of the replicator dynamic and share its local stability properties near mixed equilibria, differing mainly in how fast they converge. That reinforcement and imitation, two protocols with essentially disjoint informational requirements — one needs a rival's payoff and no own history, the other needs an own history and no rival's payoff — arrive at the same population-level equation is the strongest available evidence that the replicator equation is not an artifact of any one microfoundational story but the generic output of any adjustment process in which *something that worked keeps getting done more.* Fudenberg and Levine's comprehensive treatment situates results of this kind within a general theory in which equilibrium play emerges as the long-run outcome of a nonequilibrium learning process rather than as an a priori assumption — the individual-learning counterpart to the population-level convergence this section has just derived [20].

That convergence is also what licenses the section's economic reading. A market does not need price-taking, profit maximization, or common knowledge of the payoff structure to generate replicator-like reallocation of market share; it needs only participants who can tell, in some form, that another option is doing better than their own, and who are more likely to move toward it the larger that difference is. Where that condition fails — where switching is costly, where rivals' outcomes are unobservable, where allocation is decided by negotiated contract rather than by revision at all — the replicator equation is not a bad approximation to the market's dynamics; it is the wrong equation, and Sections 5 through 8 are about what replaces it and by how much. The punchline of this section, carried forward through the rest of the paper: the replicator equation is the dynamic of markets whose participants can see who is doing well and copy them — nothing more, nothing less.

5 The two-channel market replicator

Section 4 settled which revision protocols generate a replicator law and what a revising agent needs to know to run one. It left open who, in an actual market, is doing the revising, and against what payoff. Real markets run at least two such processes at once, on opposite sides of the same transaction: consumers who compare suppliers and switch, and firms whose capacity to serve the market expands or contracts with what they can afford to plough back into it. Nothing in the literatures surveyed so far puts these two processes into one law with separately measurable rate constants that need not even agree on which firm is winning. Doing exactly that is this paper's own contribution.

Proposition 2 (Two-channel market replicator)**.** *(a) Demand channel: consumers each attached to one supplier i; each consumer revises at rate γ by sampling another consumer and applying pairwise proportional imitation on consumer surplus u_i. Mean dynamic:*

$$\dot{s}_i = \gamma\kappa\, s_i(u_i - \bar{u}) \equiv \beta_{\mathrm{D}}\, s_i(u_i - \bar{u}).$$

(b) Supply channel (Downie–Metcalfe, attributed): firm capacity K_i grows by reinvested profit, $\dot{K}_i = r\,\pi_i K_i$ (r the retention/investment rate); with market-clearing shares $s_i = K_i/\sum_j K_j$,

$$\dot{s}_i = r\, s_i(\pi_i - \bar{\pi}) \equiv \beta_{\mathrm{S}}\, s_i(\pi_i - \bar{\pi}).$$

(c) Combined, when consumer surplus and profit are both affine in an underlying efficiency z_i ($u_i = a + bz_i$, $\pi_i = c + dz_i$ with $b, d > 0$):

$$\dot{s}_i = \beta\, s_i(z_i - \bar{z}), \qquad \beta = \beta_{\mathrm{D}}\, b + \beta_{\mathrm{S}}\, d.$$

Remark 1. The two channels are separately measurable: β_{D} from customer-flow data (churn times switching elasticity) and β_{S} from investment-cashflow sensitivities. They can disagree in sign — a firm cheap for customers but unprofitable grows through the demand channel and shrinks through the supply channel, or the reverse — and the combined β that results is exactly what Paper No. 1's selection term estimates from realized share and profitability data [101].

Part (a) is Proposition 1's mean dynamic (§4) applied to a population of consumers rather than a population of firms: a revising consumer samples another consumer, learns her realized surplus, and switches with probability proportional to the positive part of the gap — exactly Schlag's proportional-imitation rule [21], run on consumer surplus instead of firm profit. Nothing about the underlying mechanism changes; what changes is only the population doing the revising and the payoff its members compare. That the two channels below reduce to the identical mean-dynamic form is not a coincidence requiring separate justification. It is the entire content of Proposition 1: any population that revises by pairwise proportional imitation on some payoff obeys a replicator law in that payoff, whichever population and whichever payoff it happens to be.

Part (b) is not new, and the paper claims none of it as its own. Downie's neglected *Competitive Process* modeled competition as an engine of internal expansion, in which the more efficient firm earns more, retains more, and grows its productive capacity faster than its rivals — a mechanism that requires no consumer to notice or switch anything at all [48], and whose fuller history Section 2 already gives. Metcalfe made the mechanism exact by recognizing it as a market-level instance of Fisher's fundamental theorem of natural selection: the rate at which a firm's share of aggregate capacity grows is proportional to the deviation of its profitability from the population's mean profitability, which is precisely $\dot{s}_i = rs_i(\pi_i - \bar{\pi})$, so that the resulting rate of increase of mean profitability across the population is governed by the profitability variance the population currently carries [49]. His later book-length treatment extends the same argument into a full account of competition as continuous, variety-generating, and selecting rather than a state of rest reached and held [50]. What this paper adds to Downie's and Metcalfe's result is not the mechanism but the observation that it is one specific realization of the same mean-dynamic machinery driving the demand channel, differing only in which payoff is tracked and which population revises against it — capital owners reinvesting rather than consumers imitating — and that the two can therefore be measured, and can disagree, on the same footing.

A three-firm worked example

Consider three firms with market shares $s_1 = 0.40$, $s_2 = 0.35$, $s_3 = 0.25$, consumer surplus $u_1 = 9$, $u_2 = 7$, $u_3 = 5$ (dollars of surplus per unit purchased), and profit rate $\pi_1 = 0.08$, $\pi_2 = 0.12$, $\pi_3 = 0.06$ (return on invested capital). Firm 1 is the customer favorite; firm 2, less loved, is the best run.

Share-weighted mean surplus is $\bar{u} = 0.40(9) + 0.35(7) + 0.25(5) = 3.60 + 2.45 + 1.25 = 7.30$, so the demand-channel advantages are $u_1 - \bar{u} = 1.70$, $u_2 - \bar{u} = -0.30$, and $u_3 - \bar{u} = -2.30$ (weighted sum $0.40(1.70) + 0.35(-0.30) + 0.25(-2.30) = 0.68 - 0.105 - 0.575 = 0.000$, as it must be for any deviation from a share-weighted mean). Share-weighted mean profit is $\bar{\pi} = 0.40(0.08) + 0.35(0.12) + 0.25(0.06) = 0.032 + 0.042 + 0.015 = 0.089$, so the supply-channel advantages are $\pi_1 - \bar{\pi} = -0.009$, $\pi_2 - \bar{\pi} = 0.031$, and $\pi_3 - \bar{\pi} = -0.029$ (weighted sum $0.40(-0.009) + 0.35(0.031) + 0.25(-0.029) = -0.0036 + 0.01085 - 0.00725 = 0.00000$).

The two channels already disagree on firms 1 and 2 exactly as Remark 1 says they can: firm 1 carries by far the largest demand-channel advantage but a small, slightly negative supply-channel one, while firm 2 carries a negative demand-channel advantage and by far the largest supply-channel advantage. Which firm gains share is not obvious from either number alone.

Take $\beta_{\mathrm{D}} = 0.10$ and $\beta_{\mathrm{S}} = 0.20$ per revision period and set $\Delta t = 1$. The demand-channel share change $\Delta s_i^D = \beta_{\mathrm{D}}\, s_i(u_i - \bar{u})$ is

$$\Delta s_1^D = 0.10(0.40)(1.70) = 0.0680, \quad \Delta s_2^D = 0.10(0.35)(-0.30) = -0.0105, \quad \Delta s_3^D = 0.10(0.25)(-2.30) = -0.057$$

summing to $0.0680 - 0.0105 - 0.0575 = 0.0000$, confirming the demand channel alone only reallocates share. The supply-channel share change $\Delta s_i^S = \beta_{\mathrm{S}}\, s_i(\pi_i - \bar{\pi})$ is

$$\begin{aligned}
\Delta s_1^S &= 0.20(0.40)(-0.009) = -0.00072, \\
\Delta s_2^S &= 0.20(0.35)(0.031) = 0.00217, \\
\Delta s_3^S &= 0.20(0.25)(-0.029) = -0.00145,
\end{aligned}$$

summing to $-0.00072 + 0.00217 - 0.00145 = 0.00000$. Adding both channels gives the total one-step share change $\Delta s_i = \Delta s_i^D + \Delta s_i^S$:

$$\Delta s_1 = 0.0680 - 0.00072 = 0.06728, \quad \Delta s_2 = -0.0105 + 0.00217 = -0.00833, \quad \Delta s_3 = -0.0575 - 0.00145 = -0.058$$

which sums to $0.06728 - 0.00833 - 0.05895 = 0.00000$, so the post-revision shares $s_1' = 0.46728$, $s_2' = 0.34167$, $s_3' = 0.19105$ still sum to 1.00000.

Firm 1, the customer favorite, gains just under 6.7 points of share net, because its demand-channel advantage dwarfs its tiny profit disadvantage at these rate constants. Firm 2, the actually most profitable firm, still loses share net — but only 0.83 points, a small number that is small precisely because the two channels pull in opposite directions rather than because either channel is weak on its own: firm 2 loses just over one share point through the demand channel and recovers about a fifth of a point through the supply channel. Firm 3, behind on both counts, loses the most. Nothing in the mathematics fixes that ranking: raise β_{S} relative to β_{D} — a market with fast internal-finance-driven capacity growth and slow, costly consumer switching — and firm 2 can gain share outright despite being the firm customers like less. The two-channel replicator is exactly precise enough, and no more, to turn that comparative-statics question into an arithmetic one once β_{D} and β_{S} are measured separately.

The demand channel of Proposition 2 assumes every consumer switches the instant a rival offers even an infinitesimally higher surplus. Real consumers face a cost of switching — a cancellation fee, a data-export problem, a new interface to learn — and Proposition 3 adds exactly that friction to the mean dynamic derived above.

Proposition 3 (Switching-cost damping)**.** *Add a switching cost $c \geq 0$ to the demand channel: the revising consumer switches only if the sampled alternative's surplus advantage exceeds c, with probability $\kappa\,[u_j - u_i - c]_+$. Mean dynamic:*

$$\dot{s}_i = \gamma\kappa \sum_j s_i s_j \big([u_i - u_j - c]_+ - [u_j - u_i - c]_+\big).$$

For the two-firm case with surplus gap $\Delta = u_1 - u_2 > 0$:

$$\dot{s}_1 = \gamma\kappa\, s_1 s_2\, [\Delta - c]_+,$$

so selection stops entirely for $\Delta \leq c$ (a selection shadow*: quality differences smaller than the switching cost are evolutionarily invisible). With consumer-level match-value heterogeneity, the damping becomes smooth:*

$$\dot{s}_1 = \beta_{\text{D}}(c)\, s_1 s_2\, \Delta, \qquad \beta_{\text{D}}(c) = \gamma\kappa\ \Pr[G > c],$$

where G is the difference between two sampled consumers' idiosyncratic match values for the two firms. Appendix A gives the first-order derivation.

Corollary 1. *Measured selection intensity is an institutional variable: anything that raises c — contract lock-ins, data non-portability, learning costs — lowers β_{D} linearly through the tail probability of the taste-gap distribution.*

The kinked law states the sharp case bluntly: below the threshold, not one consumer, however carefully she is sampled, ever finds it worth switching, and the demand channel goes numerically silent even though firm 1 is strictly better for every consumer who samples it. This is the paper's central empirical claim in miniature: a quality gap can be real, persistent, and completely invisible to a switching-cost-weighted market, not because consumers fail to notice it but because noticing it is not enough to make acting on it worth the cost. Figure 1 plots the two-firm share path for a fixed surplus advantage $\Delta = 0.30$ under three regimes side by side: the undamped replicator, which reaches the better firm's eventual monopoly at the textbook logistic rate; the kinked law at a switching cost of $c = 0.15$ below the gap, which reaches the same eventual outcome at roughly half the rate; and the kinked law at $c = 0.35 > \Delta$, where the share path is flat for as long as the gap and the cost both hold — the selection shadow made literal, not a slower version of selection but its complete absence.

The kink itself is an artifact of assuming every consumer has identical tastes for the two firms net of u_i. Real consumers differ idiosyncratically in which supplier suits them, and that heterogeneity smooths the kink into the linear law $\beta_{\mathrm{D}}(c) = \gamma\kappa \Pr[G > c]$ of Proposition 3, plotted for Gaussian and Laplace taste-gap distributions in Figure 2 (§8); Figure 1's own third path, at the same $c = 0.15$ but with Gaussian heterogeneity of spread $\tau = 0.2$, recovers gradual selection at the damped intensity rather than the sharp kink. $\beta_{\mathrm{D}}(c)$ is monotone decreasing in c by construction — $\Pr[G > c]$ is a survival function — so raising the switching cost lowers measured demand-side selection by exactly the fraction of consumer pairs whose idiosyncratic gap no longer clears the cost, never discontinuously except in the zero-heterogeneity limit that produces the kink. Klemperer's original analysis of markets with consumer switching costs already showed that once any cost of this kind exists, ex ante identical products become ex post differentiated and current market share becomes a state variable for future profit [81, 82]; Proposition 3 gives that qualitative claim its selection-dynamics content, and Farrell and Klemperer's survey of coordination and lock-in catalogs exactly the institutional levers — contract terms, data portability, compatibility choices — that set c in practice [83].

Corollary 1 is therefore not a footnote. It says the measured selection intensity of a market is not a technological constant of that market's products; it is a policy-sensitive institutional parameter, and the switching-cost half of it is set largely by contract design, interoperability rules, and how expensive a market makes it to leave. Section 8's prediction table and Section 7's reading of the weak, damped selection coefficients recovered from panel regressions both rest on exactly this mechanism: what looks like weak selection in aggregate data can be, firm by firm, a fully switched-off channel sitting inside a shadow that no consumer needs to be irrational to remain in.

6 What evolution selects: the Walrasian surprise

Sections 5–4 established what a market replicator looks like and how it can be damped; this section asks a sharper question. Suppose a market really does run on the imitation protocol of Table 1 – firms copy whoever is currently most profitable, and only rarely experiment with something else. What does that process actually select for? The naive guess, inherited from the Alchian–Friedman "as-if" tradition (§2), is that imitation of the successful should converge on individually profit-maximizing, Cournot–Nash behavior: copy the best performer often enough and the population should end up doing what a rational optimizer would have done anyway. This guess is wrong, and the way it is wrong is the paper's title.

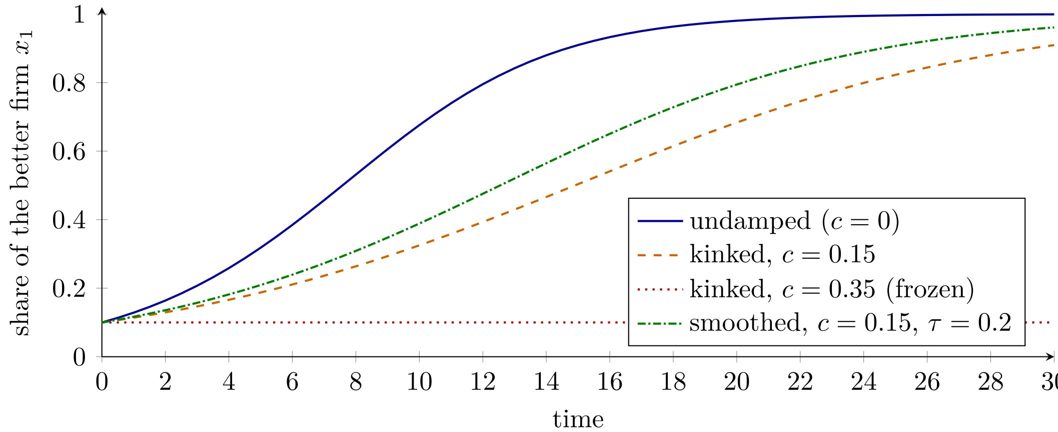

Figure 1: Two-firm share dynamics with a fixed surplus advantage $\Delta = 0.30$ for the better firm. Without switching costs the share follows the logistic replicator path. A switching cost of $c = 0.15$ under the kinked law of Proposition 3 halves the effective rate; at $c = 0.35 > \Delta$ the advantage falls inside the selection shadow and the share never moves. The smoothed law with Gaussian taste heterogeneity ($\tau = 0.2$) recovers gradual selection at the same $c = 0.15$, at the damped intensity $\beta_{\mathrm{D}}(c) = \gamma\kappa \Pr[G > c]$.

Proposition 4 (Walrasian selection)**.** *In a symmetric n-firm Cournot market where firms imitate the output of the highest-profit rival and rarely experiment, the unique long-run stochastically stable state is the Walrasian output (price equals marginal cost), not Cournot–Nash [32–34].*

The result is Vega-Redondo's, and it is worth stating in its own formal terms rather than only in the compressed language of Proposition 4, because the mechanism lives in the model's details [32].

Theorem 6.1 (Vega-Redondo, 1997)**.** *Consider n symmetric quantity-setting oligopolists with common cost function $C(\cdot)$ facing inverse demand $P(\cdot)$. Let the Walrasian output $\hat{q}$ satisfy*

$$P(n\hat{q})\,\hat{q} - C(\hat{q}) \;\geq\; P(n\hat{q})\,q - C(q), \qquad \forall\, q \geq 0,$$

i.e. $\hat{q}$ is the output a firm would choose if it took the market-clearing price as parametrically given, ignoring its own effect on price – in contrast to a Cournot–Nash firm, which internalizes that effect. Each period, each firm independently revises its output with probability $p > 0$, imitating whichever output achieved the highest profit in the previous period; independently, with probability $\epsilon > 0$ each firm instead experiments with an arbitrary output. Let $\omega_{\hat{q}} = (\hat{q}, \hat{q}, \ldots, \hat{q})$ denote the monomorphic Walrasian state, and let μ^* be the limit, as $\epsilon \to 0$, of the unique invariant distribution of the resulting Markov process. Then $\mu^*(\omega_{\hat{q}}) = 1$.

That is: as experimentation becomes vanishingly rare, all of the process's long-run probability mass concentrates on every firm playing the Walrasian output, and none of it survives on the Cournot–Nash output – even though no firm in the model is a price-taker, the population is finite, and no firm is maximizing anything beyond "copy whoever made the most money last period." The theorem's force is that it identifies a *unique* stochastically stable state out of a continuum of candidates, and that state is competitive rather than oligopolistic.

The mechanism: spite. The reason imitation overshoots Cournot–Nash and lands on the Walrasian output is that the imitation protocol of Proposition 1 does not select on the level of π_i; it selects on the sign and magnitude of $\pi_i - \bar{\pi}$, because a revising firm switches to a rival's output with probability proportional to $[\pi_j - \pi_i]_+$. Schaffer showed that maximizing relative payoff is not the same problem as maximizing absolute payoff whenever one firm's action affects a rival's payoff, and that the finite-population analogue of an ESS is generally *spiteful*: it favors any deviation that widens the profit gap in the deviant's favor, even at a cost to the deviant's own absolute profit [33]. In Cournot competition this has a concrete shape. A firm that expands output beyond the individually optimal (Cournot–Nash) level depresses the market-clearing price for every firm simultaneously, itself included; a non-strategic, atomistic optimizer would never do this, because the price externality is a pure cost with no offsetting benefit to the firm imposing it. But under imitate-the-best-profit revision, what is being selected is not the level of the deviant's profit but its profit *relative to the rest of the population*, and expansion imposes the price externality on rivals' unchanged output while the deviant captures a larger physical share of the shrunken pie. Output expansion undercuts rivals in exactly this relative sense: it can raise $\pi_i - \bar{\pi}$ while lowering π_i itself. Schaffer's independent critique of the Friedman conjecture that market selection favors profit-maximizing firms is the general statement of this point – once firms have market power, selection can favor strategies that a profit-maximizer would never choose [34]. Successive rounds of imitation propagate whichever output currently has the highest relative profit, pushing the common output upward past Cournot–Nash, until the rent that spiteful expansion was competing away has been driven to zero – which happens exactly at the Walrasian output, where price equals marginal cost and no further expansion improves relative standing because there is no markup left to redistribute. Cournot oligopoly is itself a canonical instance of a potential game, a class within which several revision-protocol dynamics are known to converge to Nash equilibrium by construction [28], which sharpens why Proposition 4's result — convergence to the non-Nash Walrasian point instead — is surprising rather than a generic feature of the underlying game's structure.

Generalizing beyond Cournot. Vega-Redondo's result is stated for a specific class of models: symmetric firms, a common cost function, and quantity competition in a homogeneous good. Alós-Ferrer and Ania show the mechanism is not an artifact of that specific setup. They establish evolutionary stability of price-taking ("act as if negligible") behavior across the broader class of symmetric, aggregative games under supermodularity – games in which each firm's payoff depends on rivals only through an aggregate statistic of their choices, of which Cournot oligopoly is one instance among many [35]. The generalization matters for how far Proposition 4 can travel: it extends the Walrasian-selection result beyond one cost function and one demand system, but it does not relax the requirements of symmetry or aggregative structure, and it retains the same imitation-plus-rare-experimentation revision protocol.

The experimental scorecard. Whether real imitating subjects actually reach the Walrasian output, rather than merely tending toward it, is an empirical question with a more qualified answer than the theorem suggests. Huck, Normann, and Oechssler ran Cournot oligopoly experiments and found that when subjects are shown rivals' quantities and profits, imitation of successful behavior does emerge and does push markets toward more competitive, higher-output, lower-price outcomes than Cournot–Nash predicts – but the convergence is directional, not exact [37]. Offerman, Potters, and Sonnemans directly compare imitation-based and belief-based (Cournot best-response) learning models against oligopoly experimental data and find that neither model alone organizes the data cleanly; observed behavior sits between the two protocols' predictions [38]. Apesteguia, Huck, and Oechssler's general theory-and-experiment treatment of imitation traces much of this disagreement to a source Table 1 already flags as decisive: differing experimental results mostly reflect differing assumptions about what information imitating subjects can actually observe about their rivals, not differing behavioral rules [36]. The honest summary is that imitation experiments confirm the *direction* of Proposition 4 – imitation pushes oligopoly play toward the competitive outcome – without confirming that real subjects land cleanly on it, and the gap tracks observability exactly as the taxonomy of §8 would predict.

Competition computes – and its precise scope. The result earns the paper's title: a population of firms, none of them solving for the competitive equilibrium, none of them even aware that a competitive equilibrium exists, can nonetheless be driven to it by a decentralized copying rule acting only on relative, realized, already-known performance. No participant is maximizing toward the Walrasian outcome; the market computes it anyway, as the stochastically stable rest point of a process whose only inputs are "who is doing best right now" and "copy them." But the theorem's scope is exactly as wide as its hypotheses and no wider. It requires symmetric firms competing in a genuinely homogeneous good, so that a single market-clearing price and a single notion of "the" Walrasian output are well defined; it requires the imitate-the-highest-profit protocol of §4 rather than a belief-based or best-response protocol, since payoff observability is what lets relative-fitness selection operate at all; and it is a statement about the limit as experimentation becomes rare, not a claim that any fixed, realistic mutation rate delivers the Walrasian state exactly. Outside that scope – heterogeneous costs, differentiated products, firms that observe rivals' strategies but not their profits, or markets where best-response rather than imitation is the operative protocol – neither the theorem nor its generalization applies, and §9 and §8 take up what happens there. Within it, the honest claim is narrow and still striking: competition, implemented as imitation of the successful, is itself a computational procedure for finding the price-taking outcome, with no optimizing agent inside it.

7 The empirical audit

Every proposition proved so far is a theorem about a stated dynamical assumption; none of it guarantees that any real market runs close to the assumption. Equation 1 makes an unusually sharp prediction once it is divided through by the share it governs: $\dot{x}_i/x_i = \pi_i - \bar{\pi}$, and under the two-channel combination of Proposition 2(c), $\dot{s}_i/s_i = \beta\,(z_i - \bar{z})$. A firm's share growth rate should be linear in its fitness measured relative to the population mean, with a slope equal to the selection intensity β and no other term. That is not a metaphorical claim; it is a regression specification, and the industrial-dynamics literature has been estimating exactly that regression, under names like "Gibrat's Law violations" and "the growth-productivity link," since long before anyone wrote it as a replicator equation [61]. What that literature finds is neither the theorem's zero (no selection) nor its textbook unit slope (frictionless selection): it finds a small, positive, remarkably stable coefficient, sitting inside a total variance of growth rates that the fitness measure explains only a modest fraction of. A parallel productivity-decomposition tradition measures the same underlying selection channel through the covariance between market share and productivity levels rather than through a growth-rate regression directly, tracing from Olley and Pakes's original telecommunications-equipment decomposition through its extension to aggregate manufacturing reallocation and net entry, its later refinement to properly attribute entry and exit's own contribution, and its application to other countries and periods, including Germany's post-reunification manufacturing sector and a cross-country comparison finding the same covariance term markedly larger in the United States than in Western Europe [66, 68–71].

7.1 Weak but positive: what the coefficient looks like

Dosi, Moschella, Pugliese, and Tamagni's four-country panel puts a number on the coefficient directly: correlated-random-effects regressions of firm growth on productivity levels and changes return elasticities with "a median across sectors of about 0.2 in all the countries" studied (France, Germany, the UK, and the US), and the regressors carrying that coefficient account for only "1/5 to 1/6 of the variance in firms' growth rates" once sector and firm heterogeneity are partialled out [64]. Bottazzi, Dosi, Jacoby, Secchi, and Tamagni's Italian and French panels sharpen the same picture at the sector level: in their worked example (sector 151), a model explaining 21.85 percent of total growth-rate variance attributes less than 5 percent of that explained variance to productivity, which works out to roughly 1 percent of the total variance in growth rates, and they report the productivity contribution as "typically below 5% and most often below 3%" across sectors, with profitability doing somewhat better at "about 5%" [65]. Read against $\dot{s}_i/s_i = \beta(z_i - \bar{z})$, this is the honest empirical content of the replicator hypothesis for real industries: the sign is right, the coefficient is detectably positive and reasonably stable across countries and time, and it is nowhere near the coefficient a frictionless, fully observable, costlessly imitated market would produce. Selection is working. It is working at a small fraction of its textbook intensity.

7.2 Selection on profit, not on physical efficiency

Foster, Haltiwanger, and Syverson's plant-level decomposition adds a qualification the growth regressions alone do not show: the fitness the market actually selects on is closer to π_i than to physical efficiency. Using output quantities rather than only revenue, they find more dispersion in physical total factor productivity than in the revenue-based measure that approximates profitability—standard deviations of 0.26 log points for physical TFP versus 0.22 for revenue TFP and 0.21 for the traditional (undecomposed) TFP measure, on a sample of 17,699 plant-years—yet survival tracks the profitability side: "a one standard deviation increase in demand shocks accounts for a decrease in the likelihood of exit that is 3 to 4 times larger than the decrease in the likelihood of exit from a one standard deviation in physical productivity" [67]. Plants exit or survive by what they earn, not directly by what they physically produce per unit of input. This is precisely the case Proposition 2(c)'s affine-in-z assumption sets aside for tractability: it posits u_i and π_i both increasing in one underlying efficiency z_i, but a plant can be physically inefficient and commercially protected (a favorable local demand shock, a differentiated variety) or physically excellent and commercially exposed. Where the two diverge, the market replicates on π_i, and a regression that measures selection against physical efficiency alone will understate β for reasons that have nothing to do with switching costs or observability—it has mismeasured the fitness variable, not failed to find selection. The dispersion this decomposition works with is not particular to one dataset or period: Syverson's broader survey of within- and across-producer productivity differences finds large, persistent gaps to be a pervasive feature of measured productivity data generally [73], and his own earlier study of US manufacturing attributes part of that persistence, even within narrowly defined four-digit industries, to demand-side product substitutability rather than technology alone—a further reason apparent selection can look weak without the underlying market having stopped selecting on anything [72].

7.3 A weakening gradient

Decker, Haltiwanger, Jarmin, and Miranda's later work turns the same regression into a time series and finds the coefficient itself declining. Comparing the 1990s to the post-2000 period, they report that "the marginal responsiveness of employment growth to business-level productivity has weakened," with the post-2000 responsiveness for young high-tech firms running "only about half (in manufacturing) to two thirds (economy wide) of the peak in the 1990s" [75]. Read through this paper's mapping, that is a direct statement about β: not a claim that fitness variance shrank—the dispersion of the shocks hitting individual businesses did not fall over the same window—but a claim that the rate at which share responds to a unit of fitness advantage fell, by roughly a third to a half, in one generation. The productivity consequence is not academic: the same study estimates that the elevated responsiveness of the 1980s and 1990s contributed "as much as half a log point" of annual industry-level TFP growth in high-tech manufacturing by the mid-1990s, and that its subsequent decline cost "as much as a two-log-point drag" on annual industry TFP by 2010 [75]. A market can carry exactly the fitness dispersion a replicator needs and still replicate more slowly than it once did, if something in its revision protocol—observability, switching cost, the reinvestment coupling of §5—has changed. The measured β is not a constant of nature; this is an economy-scale instance of the damping this paper's own Proposition 3 predicts should be possible. A companion strand of the same literature documents a parallel decline on the entry margin itself: Decker, Haltiwanger, Jarmin, and Miranda separately find the US business startup rate and the share of employment at young, dynamic firms have both fallen for decades, evidence that the entry-driven component of the economy's overall selection pressure has weakened alongside the responsiveness component this section has just measured [74].

7.4 Heavy tails: lumpy revision, not smooth flow

A second, independent line of evidence concerns not the mean of the growth-rate-versus-fitness relationship but its distribution. If revision opportunities arrived as a dense stream of small, independent events—the continuous-time idealization behind Equation 1 and Proposition 1's mean dynamic—the central limit theorem would push realized growth rates toward something close to Gaussian around the deterministic path. That is not what plant- and firm-level data show. Stanley et al., using the full population of publicly traded US manufacturing firms from 1975 to 1991, find that "the distribution of annual (logarithmic) growth rates has an exponential form," with the width of the distribution shrinking as sales grow according to a power law with an exponent of approximately -0.2 across several size measures [62]. Bottazzi and Secchi confirm the same tent-shaped, Laplace rather than Gaussian, growth-rate law on independently disaggregated data, and build a model in which each firm's ability to exploit new opportunities compounds on the opportunities it has already exploited [63]. The size distribution these growth-rate laws generate is itself a well-documented power law close to the

Zipf case, exponent near one, in the full population of US firms [92], and Gabaix's survey of power-law regularities in economics and finance situates that finding within the broader class of proportional random-growth mechanisms known to generate it [93]. A Laplace distribution has exponential rather than Gaussian tails: big jumps are far more common, relative to the bulk of small changes, than a smooth aggregation of many independent micro-revisions would produce. That is consistent with the revision events that actually move a firm's share being lumpy and occasional—an acquisition, a plant closing, a large contract won or lost—rather than the continuum of small pairwise imitation events Proposition 1's mean-field derivation idealizes, and it is a second, distribution-shaped reason the same panels return small average coefficients: a regression run on annual growth rates is averaging mostly-quiet years against a minority of years carrying most of the actual selection.

7.5 The diffusion evidence: Griliches's logistic as replicator

The oldest quantitative study in this tradition already is a two-strategy replicator estimated before the word existed. Griliches modeled the state-by-state adoption of hybrid corn seed in the United States as a logistic, S-shaped curve over time [77]; the per-state parameters he fit are not recoverable from the surviving accessible record, but the functional form is exactly what Equation 1 predicts for two strategies with a constant fitness gap. Let x be the adopting share and $1 - x$ the non-adopting share, with $\pi_1 - \pi_2 = r$ constant:

$$\dot{x} = x(\pi_1 - \bar{\pi}) = x(1 - x)(\pi_1 - \pi_2) = r\,x(1 - x),$$

which is the logistic equation, whose solution is the logistic diffusion curve. A constant yield and profit advantage for hybrid over open-pollinated seed, combined with farmers who observe and copy their better-performing neighbors—pairwise proportional imitation in an agricultural setting—is sufficient to generate the S-curve Griliches measured; the state-by-state differences in timing and steepness he documented are, on this reading, differences in r and in the rate at which farmers sampled and copied one another, not evidence against a replicator process. Bass's diffusion model, fit to eleven consumer durable goods with separate coefficients for innovation (adoption independent of others) and imitation (adoption driven by prior adopters), generalizes the same S-curve family to settings where an external adoption pressure augments the peer-copying term [78]. Rogers's canonical synthesis of diffusion-of-innovation theory—adopter categories, the S-shaped diffusion curve, communication- channel effects—is the field's standard reference point for the same shape of result, and it complements rather than competes with the Griliches and Bass curves just discussed [91]. Young formally decomposes the mechanisms that can produce an S-curve—contagion, social influence, and social learning—and re-examines the historical hybrid-corn adoption data directly, showing more than one of these mechanisms is consistent with the same aggregate curve [79]. Silverberg, Dosi, and Orsenigo's disequilibrium simulation model shows the same quasi-logistic diffusion and shifting market shares emerging directly from population dynamics among diverse, boundedly rational firms, an explicit replicator-style treatment of diffusion built independently of the Griliches–Bass tradition [51]. Geroski's survey of the wider technology- diffusion literature classifies the field into epidemic, probit, density-dependence, and information-cascade families, each generating the same qualitative S-shape from a different microfoundation [80]. Only the epidemic/contagion family is a replicator process in this paper's sense—Bass's imitation term and Griliches's peer-adoption channel belong to it; the probit family (adoption triggered independently once a threshold is crossed) and pure information cascades do not require anyone to copy a better-paid neighbor at all. The S-curve is common evidence for a family of mechanisms of which the replicator is one well-fitting member, not proof of it alone.

7.6 Verdict

Put the pieces together and they tell one consistent story rather than four separate ones. The sign of the growth-versus-fitness relationship is right everywhere it has been measured; its magnitude is a small fraction of the textbook value; the fraction is smaller still when the fitness variable used is physical efficiency rather than profit, because profit and not physical efficiency is what survival actually tracks; the fraction has been shrinking further over the last three decades in the country with the best data to measure it; and the growth-rate distributions surrounding that small mean coefficient are heavy-tailed in exactly the way lumpy, occasional revision rather than smooth continuous imitation would produce. None of this is the replicator failing an empirical test. It is the replicator succeeding at a damped intensity, and the same damping this paper's Proposition 3 derives from an institutional switching cost is a mechanism these results are independently consistent with, coefficient by coefficient. A different empirical route arrives at a comparable-magnitude shortfall from the frictionless benchmark: Hsieh and Klenow's cross-country misallocation accounting finds that reallocating capital and labor toward marginal-product equality in China and India, to the efficiency observed in US data, would raise aggregate manufacturing TFP substantially, the same qualitative gap this section's coefficients are measuring industry by industry through a growth-rate lens instead [76]. The verdict this audit supports is the one stated at the outset of this program: the market is a replicator system running at a small fraction of its textbook selection intensity, and the damping is measurable institution by institution. It is also exactly the selection term the companion paper in this series audits from the demand side, using plant-level dispersion and reallocation evidence rather than growth-rate regressions [101]: the same small, real, institution-dependent β turns up under both lenses. Appendix C's simulation makes the point precisely rather than suggestively: regressing growth rate on relative fitness recovers a pure replicator's β to within a few percent, recovers a damped fraction close to the true attenuation once a Proposition 3 switching cost is imposed, and returns a coefficient statistically indistinguishable from zero, at several times the standard error, once realistic Laplace-tailed noise alone is added with no selection at all (Figure 3). The small, noisy, positive coefficients this section has surveyed across four decades of industrial-dynamics research are what a damped replicator running inside heavy-tailed noise is supposed to produce. They are not, on inspection, a case against the replicator; they are a measurement of how much of it survives the institutions real markets actually have.

8 The institutional taxonomy

The preceding sections derived one dynamic from one protocol under one channel structure: imitate on realized payoff, and the population moves according to the replicator equation, on the demand side or the supply side or both. Real markets differ in which protocol their participants actually run, and by how much. This section states five institutional dials — observable features of a market that a regulator, a market-research analyst, or a competitor could name without reference to any of this paper's mathematics — and maps each one to the formal object, already introduced, that it sets. The result is a systematic route from a market's institutions to a prediction about its share dynamics, stated precisely enough that the prediction can fail.

Five dials

Dial 1: payoff observability. The question is whether a revising consumer or firm can see who is doing well. Table 1 in §4 makes the dial formal: if participants observe a rival's realized payoff and nothing more, the available protocols are imitative — pairwise proportional imitation [21] or imitation of success [22] — and the resulting dynamic is the replicator, exactly or up to a speed factor. If participants instead know the whole payoff function and can compute a best response, the protocol shifts to the belief-based family: best response [26] or its noisy, logit-perturbed cousin [25], and the resulting dynamic is the best-response or logit dynamic, not the replicator, with qualitatively different convergence behavior (jumps to a best response rather than gradual reweighting). Markets built around visible, comparable public signals of quality — reviews, published prices, foot traffic — sit on the imitative end of this dial; markets in which quality is privately negotiated and rivals' terms are confidential sit on the belief-based end, to whatever extent belief formation is possible at all under Sandholm's general revision-protocol accounting [10]. Management practice is one plausible candidate for the trait a revising firm would need to observe on this dial: the World Management Survey methodology shows measured management quality varies widely across firms and countries and correlates with productivity, profitability, and survival, which is what makes it a heritable, selectable trait in the first place [88], and randomized and natural-experiment evidence finds the trait itself transferable across firms through targeted consulting and training interventions [89, 90].

Dial 2: switching cost and lock-in. Proposition 3 makes this dial formal directly: a switching cost c damps the demand-side intensity from β_{D} to $\beta_{\mathrm{D}}(c) = \gamma\kappa \Pr[G > c]$, the fraction of consumer pairs whose idiosyncratic taste gap exceeds the cost. Figure 2 plots this damping ratio against c for two taste-gap distributions; in both cases a switching cost on the order of one standard deviation of match-value dispersion removes roughly half the demand-side selection intensity the market would otherwise carry. Contractual lock-in, data non-portability, and firm-imposed loyalty structures are the direct economic content of c in this dial, and the literature that measures them — Klemperer's switching-cost equilibria [81, 82] and Farrell and Klemperer's survey of coordination and lock-in [83] — is where an applied estimate of c, and hence of the damping this paper's Proposition 3 predicts, would have to start.

Dial 3: payoff externalities. When a participant's payoff depends not only on its own strategy but on the population's current composition — network effects, in the language of that literature — fitness becomes frequency-dependent, $\pi_i(x)$ rather than a fixed π_i. The replicator's functional form survives this substitution unchanged, $\dot{x}_i = x_i\,(\pi_i(x) - \bar{\pi}(x))$, but its qualitative behavior does not: frequency-dependent payoffs of this kind routinely generate multiple evolutionarily stable states and history-dependent lock-in onto one of them, exactly the mechanism Katz and Shapiro and Arthur formalized for standards competition and increasing-returns adoption [84, 85], with the applied strategy literature built on top of it [83, 86]. This paper does not develop the multiple-equilibrium case; it is the subject of a later paper in this series, and dial 3 is flagged here only to locate where that territory begins.

Dial 4: reinvestment coupling. This is the taxonomy's surprising dial, and it earns its own paragraph because the intuitive prediction is wrong. The supply-side channel of Proposition 2(b) sets $\dot{K}_i = r\,\pi_i K_i$: capacity grows in proportion to reinvested profit, so β_{S} is large exactly when a firm's own realized profit is the thing financing its own growth — Downie's and Metcalfe's mechanism [48–50]. The intuitive guess is that deeper, more developed capital markets make this channel stronger, by giving good firms more ways to raise money. The mechanism runs the other way. Once external finance is frictionless and forward-looking, a firm's growth is financed against its *expected* future profitability, assessed by lenders and investors who do not need to wait for retained earnings to accumulate — exactly the entry and investment margin that Hopenhayn's, Jovanovic's, and Ericson and Pakes's industry-dynamics models put at the center of firm growth [95–97], and that Luttmer's calibration credits with roughly half of aggregate output growth once selection and entrant imitation are both accounted for [94]. When financing decouples from current realized profit in this way, β_{S} shrinks toward zero: a firm can grow while π_i is negative, or fail to grow while π_i is high, because the capital market is pricing tomorrow's profit rather than metering today's. The paradox is exactly this: financial development can make a market *less* evolutionary in the backward-looking, Downie–Metcalfe sense of the word, while making it simultaneously more forward-looking, in the Schumpeterian sense of entry and investment tracking expected rather than realized fitness. A market with a shallow banking sector and internally financed firms is, on this dial alone, the more literal replicator; a market with deep venture and public capital markets is not, even though it may reallocate capital toward good ideas faster by an entirely different, non-replicator route.

Dial 5: match heterogeneity. The dispersion of idiosyncratic consumer-firm taste, the density g behind the symmetric gap G in Proposition 3, does two things at once: it sets the tail probability $\Pr[G > c]$ that determines how much of the demand-side intensity a given switching cost damps, and, independently of any switching cost at all, it sets a ceiling on how much share the objectively best product can win, because consumers with a strong enough idiosyncratic taste for a worse alternative will not switch even absent any cost to doing so. A market in which taste is nearly homogeneous behaves close to the two-firm limit of Proposition 3, winner-take-most; a market with wide match-value dispersion caps every firm's achievable share well short of one, and can sustain persistent multi-firm coexistence with no appeal to product differentiation on the supply side at all. Kirman's model of ant-like recruitment between two objectively identical food sources is the cleanest illustration of what this dial looks like at its limit of zero underlying quality gap: pure imitation dynamics still generate switching and herding, but with no fitness content whatsoever, a case the taxonomy's other four dials must be read against as a floor rather than a signal [100].

A prediction table

Table 2 states, for a small set of recognizable market types, where the five dials plausibly locate them and what that location predicts about the statistical signature of their share dynamics. The table is deliberately built out of testable claims about the autocorrelation and fitness-sensitivity of share changes, not assertions of established fact: no claim in it has been estimated in this paper, and each row is falsifiable independently of the others.

Measurement recipes

Two of the taxonomy's objects, β_{D} and β_{S}, are not just theoretical placeholders; each has a natural empirical recipe, though neither has, to our knowledge, been estimated in exactly this

Market type	Dominant dial setting	Predicted position	Testable signature (hedged)
Franchised food service	high observability; low c (dial 1, 2)	most replicator-like	share changes should correlate with lagged relative quality/price measures, with little unexplained autocorrelation left over
Retail chains (grocery, apparel)	high observability; low c (dial 1, 2)	near-replicator; fastest convergence	quarterly share changes should track relative price and assortment fitness with a stable, estimable $\hat{\beta}$ across periods
Consumer apps with portable accounts	high observability; near-zero c; but network externalities present (dial 1, 2, 3)	replicator-like on average, punctuated	share gains should be fitness-sensitive in the typical period but exhibit occasional large discontinuous jumps consistent with tipping, not pure smooth reallocation
Enterprise software (multi-year contracts)	low pre-contract observability; high c; often auction allocation (dial 1, 2)	least replicator-like	share changes should show weak correlation with contemporaneous fitness and instead cluster at contract-renewal dates
Banking relationships (retail, commercial)	high c; low match heterogeneity relative to switching friction (dial 2, 5)	least replicator-like	share changes attributable to switching should be rare, and, conditional on occurring, insensitive in size to small fitness gaps — a detectable selection shadow
Aerospace and defense procurement	discrete auction allocation; very low revision frequency (dial 1, 2)	least replicator-like	market share should move almost only at contract-award dates, with near-zero correlation between interim performance and interim share

Table 2: Predicted position on the replicator-likeness spectrum for six market types, read off the five institutional dials. Each cell states a hedged, testable claim about the autocorrelation or fitness-sensitivity of measured share changes, not a documented finding.

form, and both recipes are stated here as proposals rather than as reports of existing estimates.

The demand-side recipe follows directly from Proposition 3's construction. β_{D} is a rate at which consumers who sample a better alternative actually switch to it; its observable components are a churn rate (how often consumers revise at all, the empirical counterpart of γ) and a switching elasticity (how much of a realized surplus advantage a revising consumer requires before switching, the empirical counterpart of κ and, once a cost is present, of $\Pr[G > c]$ itself). The natural design pairs a consumer panel that records both realized supplier choice and a measure of relative surplus (price, quality rating, or a revealed-preference proxy) with a source of variation in the switching cost itself — a regulatory change that lowers c directly, such as number portability or data-portability mandates, of the kind the switching-cost literature treats as its central comparative static [81–83]. Estimating churn's sensitivity to relative surplus before and after such a change recovers $\beta_{\mathrm{D}}(c)$ at two points on Figure 2's curve directly, and a specification that nets out imitation-without-a-fitness-gradient, in the spirit of Kirman's

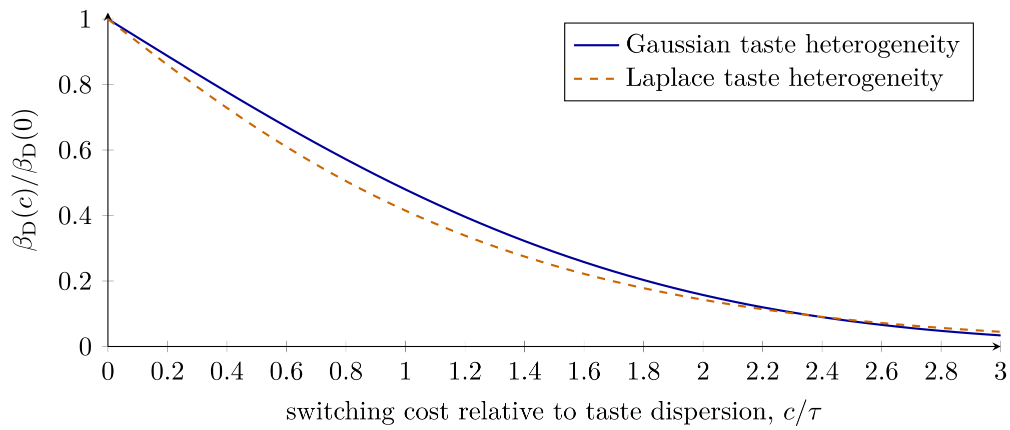

Figure 2: The damping factor of Proposition 3, $\beta_{\mathrm{D}}(c)/\beta_{\mathrm{D}}(0) = 2\Pr[G > c]$, as a function of the switching cost measured in units of the taste dispersion τ. G is the difference of two independent match values. Heavier-tailed taste heterogeneity (Laplace) preserves more selection at high switching costs because more consumer pairs sit far apart; under either distribution a switching cost equal to one standard deviation of taste removes roughly half of the demand-side selection intensity.

ant model [100], guards against attributing pure herding to selection.

The supply-side recipe is the investment-cashflow analogue: β_{S} is the sensitivity of a firm's capacity or capital growth to its own lagged profit rate, holding the market-clearing share definition of Proposition 2(b) fixed. Where Downie's and Metcalfe's mechanism is operative, this regression coefficient *is* β_{S}, directly [48, 49]; where dial 4's paradox is operative, the same regression, run on firms with ready access to external, forward-looking finance, should recover a coefficient attenuated toward zero relative to internally financed firms in the same industry, holding profitability dispersion fixed. Firm-level panels of the kind the industrial-dynamics literature already builds for productivity and growth regressions [64, 67] are the right shape of data for this recipe; what they do not yet report is the split, by financing regime, that dial 4 predicts should matter.

Both recipes feed the same downstream object. Once β_{D} and β_{S} are separately estimated, Proposition 2(c)'s combined intensity $\beta = \beta_{\mathrm{D}} b + \beta_{\mathrm{S}} d$ is exactly the selection-term coefficient this paper's companion treats as an estimated, market-structure-dependent input to its own accounting identity [101] — which is where §10 picks the thread back up.

9 Where the analogy breaks

Sections 5–8 treated the replicator equation as a dial-adjustable description of market competition: strong where payoffs are observable and switching is cheap, damped where they are not, but everywhere still recognizably the same dynamic underneath. That treatment has limits that damping does not describe, because damping weakens a coefficient while leaving the equation's structure intact, and four features of real competition instead violate a structural assumption the equation makes outright. Each is stated here as precisely as the honest-positioning commitment of §1 requires: what fails, who showed it, and what has to replace the replicator once it does.

Foresight. The replicator's growth rate $\dot{x}_i = x_i(\pi_i - \bar{\pi})$ is a function of currently realized payoffs alone; nothing on the right-hand side depends on any belief about π_i or $\bar{\pi}$ at a future date. This is not an approximation the equation makes for convenience – it is definitional of a backward-looking revision protocol, inherited directly from Proposition 1's mean dynamic, in which a firm's propensity to be imitated depends on last period's profit and nothing else. Real entry, exit, and investment decisions are not backward-looking in this sense: a firm enters, expands capacity, or exits by comparing an expected discounted value of continuing against a cost, under uncertainty about a payoff it has not yet realized. Hopenhayn's stationary-equilibrium model of entry and exit makes the exit margin an explicit function of a firm's expected continuation value given its current productivity draw and the distribution of future shocks, not of its realized share of current industry profit [95]. Jovanovic's selection model has firms learning their own efficiency through a Bayesian updating process on noisy signals and expanding or exiting according to the posterior expected value of continuing to operate – a forward, belief-based calculation with no counterpart in a fitness-proportional share update [96]. Ericson and Pakes generalize both into a Markov-perfect framework in which entry, exit, and investment are jointly the solution to a dynamic program over an anticipated future industry state [97]. What replaces the replicator here is not a modification of its coefficients but a different object entirely: a value-function-based decision rule, in which the state variable driving a firm's growth is an expectation, not a realized relative payoff. A related countervailing case sits on the incumbent-pricing side of the same foresight problem: contestable-markets theory shows the mere threat of entry, rather than the realized turnover a replicator dynamic tracks, can discipline incumbent behavior and produce competitive-looking outcomes with no selection-through-turnover mechanism operating at all [87].

Increasing returns. Proposition 2(c) already allows fitness to depend on the population state through an underlying efficiency z_i, so ordinary frequency dependence is not itself a break. What fails is a stronger structure: when a firm's payoff rises *in its own share* – when being bigger is directly what makes a product more attractive, as under network externalities – the resulting dynamical system can acquire multiple stable states whose basins of attraction depend on where the population starts, not on which firm has the higher fitness. The folk-theorem package of §3 says stable rest points are Nash and ESS are asymptotically stable; it says nothing about which of several such stable states a given trajectory reaches, and with strong enough positive feedback that choice is no longer a fitness ranking at all but a path-dependent, near-permanent accident of history. Arthur's model of competing technologies under increasing returns to adoption shows market share evolving as a non-ergodic, history-dependent stochastic process – formally a generalized urn scheme, not a deterministic fitness gradient – in which early, arbitrarily small events can tip the market onto a technology that need not be the more efficient one, and the tip is irreversible [85]. Katz and Shapiro's model of network externalities and compatibility shows the same structure from the demand side: multiple equilibria, potential lock-in to an incompatible or inferior standard, and a strategic compatibility game layered on top of what looks superficially like ordinary competition [84]. What replaces the replicator's global, fitness-driven convergence is a non-ergodic, path-dependent stochastic process together with the compatibility/coordination game theory needed to analyze which of several lock-ins occurs – a question of history and strategic standard-setting, not of relative efficiency. This break also marks the bridge to the territory of Paper No. 7, noted only in passing here.

Innovation. The replicator equation $\dot{x}_i = x_i(\pi_i - \bar{\pi})$ is defined over a fixed, finite index set $i = 1, \ldots, n$: it can only reweight the shares of strategies that already exist in the population. There is no term on its right-hand side capable of introducing a genuinely new strategy $n + 1$ with positive share where there was none before – selection, in this equation, is pure reallocation among incumbents, never creation. Page and Nowak's unification of replicator, adaptive, and mutator dynamics into a common quotient framework is the formal statement of what must be added to escape this limit: a mutation or innovation term that continuously feeds a small, non-fitness-driven inflow of probability mass onto strategies with zero current share, turning the pure replicator into a replicator-mutator process capable of genuine novelty rather than only reallocation [14]. What this means for the market analogy is that everything in §5–§6 describes how an existing set of products, technologies, or firms gets reweighted by competition; it says nothing about where a new product, technology, or firm comes from in the first place. Working out the replicator-mutator extension for markets – how R&D, entrepreneurial entry, and product innovation inject new strategies into the simplex the replicator otherwise only reallocates within – is Paper No. 3's subject, not this one's.

The unit problem. Propositions 1–4 all presume a population of atomic, comparably scaled units, each indexed by a single i, each carrying one payoff π_i and one share x_i or s_i in one simplex. A multi-product firm or a conglomerate does not fit this indexing. A single legal firm competes simultaneously in several distinct product markets, each with its own rivals, its own demand conditions, and its own local π and $\bar{\pi}$; the firm's aggregate "market share" is a composite of shares in markets that need not move together at all – a conglomerate can be gaining share in one line while losing it in another, with no single well-defined trajectory $x_i(t)$ that represents "the firm." Applying the two-channel replicator of §5 to such a firm requires deciding, prior to any measurement, what the unit i actually is, and the propositions themselves are silent on the question because they take the indexing as given. Paper No. 1's selection-accounting framework confronts an equivalent version of this problem in its own terms, and resolves it not by insisting the firm itself is the replicating unit but by locating the unit of inheritance-like continuity in the *routine* – the standard operating procedure, pricing rule, or organizational competence that persists, gets copied across divisions, and diffuses across firms independently of any single firm's survival [101]. This is Nelson and Winter's original move: routines, not firms, are the bearers of heritable continuity, while the firm is the interactor that competes for share and is selected on the profitability its routines jointly generate [47]. Hodgson and Knudsen generalize the replicator/interactor distinction into a full ontology for exactly this purpose [55], and Knudsen shows explicitly how internal, organizational-level selection among routines can be reconciled with market-level selection among the firms that carry them, rather than treated as a competing account [59]. What replaces the single-index replicator for multi-product firms is therefore a multi-level accounting – routines or business lines as the base replicating units, firms as interactors bundling many of them – a formalization this paper does not carry out and leaves, along with the innovation break above, to the routines paper to come.

10 A measurement program

The taxonomy of §8 is a set of falsifiable rankings, not yet a set of estimates. This section states what data would turn the rankings into numbers, how the resulting estimates plug into the accounting this paper's companion built, and what the next two papers in this series add that this one deliberately leaves out.

What would estimate β_D and β_S industry by industry

The two measurement recipes of §8 point to two distinct kinds of panel. Estimating β_D industry by industry requires a consumer- or account-level panel that records realized supplier choice, a usable proxy for relative surplus, and, ideally, a dated change in the switching cost itself — a portability mandate, a contract-length regulation, a platform policy change — so that the same industry can be observed at two points on the damping curve of Figure 2 rather than at one. Estimating β_S industry by industry requires a firm-level panel with capacity or capital measures, profit rates, and a way to split firms by financing regime — bank-dependent versus public-market-financed, young versus seasoned, in a jurisdiction with shallow capital markets versus one with deep ones — of the kind already used to separate physical from revenue-based selection in the productivity literature [67]. Neither panel is exotic; what is missing is not data infrastructure but the specific cut, by financing regime and by dated switching-cost change, that the taxonomy's dials 2 and 4 say should matter and that existing growth-productivity regressions do not report.

The taxonomy's rankings are a test, not an assumption

Table 2 is falsifiable as a set, not only row by row. If franchised food service, retail chains, and portable-account consumer apps do not show larger, more fitness-correlated share movements than enterprise software, banking relationships, and aerospace procurement, once churn frequency is held constant, the taxonomy's central premise — that observability and switching cost, not industry size or growth rate, are what set a market's position on the replicator spectrum — is wrong, independent of whether any individual dial's mechanism is right in isolation. A weaker but still informative test follows dial 4 alone: an industry with two segments differentiated mainly by ease of external financing should show a measurably smaller β_S in the better-financed segment, holding profitability dispersion fixed; finding the opposite would falsify the financial-development paradox specifically, without implicating the rest of the taxonomy.

Where this plugs into the wider accounting

The combined intensity $\beta = \beta_D b + \beta_S d$ of Proposition 2(c) is not a new object invented for this paper's own purposes; it is exactly the selection-term coefficient that this paper's companion treats, in its own Fisher-theorem analogue, as an estimated, market-structure-dependent input rather than a structural constant [101]. That paper asks how much of aggregate productivity growth a measured β explains, given the cross-sectional variance of the trait it acts on; this paper supplies the institutional account of why β takes the value it does, industry by industry, before any productivity growth is measured at all. The two papers' drift benchmarks connect the same way: this paper's channels operate against a demographic noise floor, and any estimate of β_{D} or β_{S} from a real industry panel is only informative once it is checked against that floor, industry concentration included, rather than read off a bare regression coefficient.

What the rest of the series adds

Two limits of this paper are deliberate, and two later papers take them up directly rather than folding hurried extensions in here. The replicator equation, however it is derived, only reweights strategies that already exist in the population; it cannot generate a genuinely new competitor, a new format, or a new business model, because there is no mass on that strategy to select for [14]. A later paper in this series takes up the replicator-mutator extension and treats entry and innovation as the process that keeps refilling the simplex this paper's dynamics only reallocate across. And this paper's measurement recipes, once run on real panels, will recover β_{D} and β_{S} estimates whose statistical significance depends on knowing how much apparent selection a market of a given concentration would show under pure drift; a further paper in the series extends the drift-null benchmark from productivity growth, where the companion paper states it, to market-share panels specifically, so that a measured β_{D} or β_{S} can be judged against the right noise floor rather than an assumed one.

Conclusion

The argument of this paper has run in one direction throughout: from an observable market institution, to the revision protocol it implements, to the population dynamic that protocol produces, to a rate that a real panel could in principle measure. None of the individual links is new — Schlag's imitation rule, Downie's and Metcalfe's reinvestment mechanism, Vega-Redondo's Walrasian result, Sandholm's protocol framework all predate this paper and are credited to their authors throughout it. What the chain adds, once assembled and read against the industrial-dynamics literature's own measured magnitudes, is a specific answer to a question usually left at the level of metaphor: not whether competition computes, but how much computing a given market's institutions let it do. The honest verdict of §7 is that most markets run this computation at a small fraction of its textbook rate. The taxonomy of §8 says that fraction is not noise; it is set, dial by observable dial, by choices about visibility, lock-in, and finance that a market's designers already make for other reasons, and could in principle measure the evolutionary consequences of before they make them.

A Proofs

A.1 Proposition 1: pairwise proportional imitation implies the replicator equation

Recall Sandholm's general mean-dynamic equation for a revision protocol $\rho_{ij}(\pi, x)$, the rate at which an i-strategist switches to j:

$$\dot{x}_i \;=\; \sum_j x_j\,\rho_{ji}(\pi, x) \;-\; x_i \sum_j \rho_{ij}(\pi, x).$$

Under pairwise proportional imitation [21], a revising agent samples one other agent uniformly at random from the whole population, so the probability the sampled partner plays j is exactly

x_j; conditional on that meeting, the revising agent adopts j with probability $\kappa\,[\pi_j - \pi_i]_+$. The full state-dependent protocol is therefore the product of the meeting probability and the conditional switch probability,

$$\rho_{ij}(\pi, x) = x_j\,\kappa\,[\pi_j - \pi_i]_+, \qquad \rho_{ji}(\pi, x) = x_i\,\kappa\,[\pi_i - \pi_j]_+,$$

the second obtained from the first by exchanging the roles of i and j. Substituting both into the general mean dynamic:

$$\dot{x}_i = \sum_j x_j \cdot x_i\kappa[\pi_i - \pi_j]_+ \;-\; x_i \sum_j x_j\kappa[\pi_j - \pi_i]_+ = x_i\kappa \sum_j x_j \Big([\pi_i - \pi_j]_+ - [\pi_j - \pi_i]_+\Big).$$

Lemma. For any real a, $[a]_+ - [-a]_+ = a$. *Proof.* If $a \geq 0$: $[a]_+ = a$ and $[-a]_+ = 0$, so the difference is a. If $a < 0$: $[a]_+ = 0$ and $[-a]_+ = -a$, so the difference is $0 - (-a) = a$. Both cases give a. □

Applying the lemma with $a = \pi_i - \pi_j$ (so $[\pi_j - \pi_i]_+ = [-a]_+$) collapses the bracket to $\pi_i - \pi_j$ exactly, with no approximation:

$$\dot{x}_i = x_i\kappa \sum_j x_j(\pi_i - \pi_j) = x_i\kappa\Big(\pi_i \sum_j x_j - \sum_j x_j\pi_j\Big) = x_i\kappa(\pi_i - \bar{\pi}),$$

using $\sum_j x_j = 1$ and $\bar{\pi} = \sum_j x_j\pi_j$ in the last step. This is $\dot{x}_i = \kappa\,x_i(\pi_i - \bar{\pi})$, Proposition 1, with every step exact and no term dropped. □

A.2 Proposition 2(a),(b): the two channels via the quotient rule for shares

(a) Demand channel. Consumers are the revising population, consumer surplus u_i is the payoff, and the consumer population is fixed (no entry or exit, only switching), so consumer shares s_i are population shares in exactly the sense of A.1. Each consumer receives a revision opportunity at per-capita rate γ and, on revising, applies pairwise proportional imitation on u_i with intensity κ; the composite protocol is $\rho_{ij}(\pi, x) = \gamma\,s_j\,\kappa\,[u_j - u_i]_+$, identical in form to A.1 with the substitutions $x \to s$, $\pi \to u$, $\kappa \to \gamma\kappa$. The derivation of A.1 applies verbatim under this relabeling and yields $\dot{s}_i = \gamma\kappa\,s_i(u_i - \bar{u}) \equiv \beta_{\mathrm{D}}\,s_i(u_i - \bar{u})$, Proposition 2(a).

(b) Supply channel [attributed to 48, 49]. Capacity grows by reinvested profit, $\dot{K}_i = r\,\pi_i K_i$. Let $K = \sum_j K_j$ be industry capacity and $s_i = K_i/K$. First, the growth rate of the aggregate:

$$\dot{K} = \sum_j \dot{K}_j = \sum_j r\pi_j K_j = rK \sum_j \frac{K_j}{K}\pi_j = rK \sum_j s_j\pi_j = rK\bar{\pi}.$$

Now differentiate the share by the quotient rule, $\dot{s}_i = \dfrac{\dot{K}_i K - K_i\dot{K}}{K^2}$, and split the right-hand side:

$$\dot{s}_i = \frac{\dot{K}_i}{K} - \frac{K_i}{K} \cdot \frac{\dot{K}}{K} = \frac{\dot{K}_i}{K} - s_i\frac{\dot{K}}{K}.$$

Substitute $\dot{K}_i = r\pi_i K_i = r\pi_i s_i K$ and $\dot{K} = rK\bar{\pi}$ from above:

$$\dot{s}_i = \frac{r\pi_i s_i K}{K} - s_i \cdot \frac{rK\bar{\pi}}{K} = r\pi_i s_i - rs_i\bar{\pi} = r\,s_i(\pi_i - \bar{\pi}).$$

Writing $\beta_{\mathrm{S}} := r$ gives exactly $\dot{s}_i = \beta_{\mathrm{S}}\, s_i(\pi_i - \bar{\pi})$, Proposition 2(b), derived with no approximation: the quotient rule is exact, and every substitution used only the two defining equations $\dot{K}_i = r\pi_i K_i$ and $s_i = K_i/K$.

(c) The combined channel. Both dynamics act on the same share variable s_i simultaneously, so they add: $\dot{s}_i = \beta_{\mathrm{D}}\, s_i(u_i - \bar{u}) + \beta_{\mathrm{S}}\, s_i(\pi_i - \bar{\pi})$. Under $u_i = a + bz_i$, the mean cancels the intercept, $u_i - \bar{u} = b(z_i - \bar{z})$, since $\bar{u} = a + b\bar{z}$; likewise $\pi_i - \bar{\pi} = d(z_i - \bar{z})$ under $\pi_i = c + dz_i$. Substituting,

$$\dot{s}_i = \beta_{\mathrm{D}}\, s_i\, b(z_i - \bar{z}) + \beta_{\mathrm{S}}\, s_i\, d(z_i - \bar{z}) = \big(\beta_{\mathrm{D}} b + \beta_{\mathrm{S}} d\big)\, s_i(z_i - \bar{z}) = \beta\, s_i(z_i - \bar{z}),$$

with $\beta := \beta_{\mathrm{D}} b + \beta_{\mathrm{S}} d$, exactly Proposition 2(c). $\square$

A.3 Proposition 3: the kink, and its first-order smoothing under taste heterogeneity

The kinked law. Adding a switching cost $c \geq 0$ changes the conditional switch probability in A.2(a) to $\kappa[u_j - u_i - c]_+$, so the protocol becomes $\rho_{ij}(\pi, x) = \gamma\, s_j\, \kappa\, [u_j - u_i - c]_+$. Substituting into the general mean dynamic exactly as in A.1, but *without* invoking the lemma (the "$-c$" shift breaks the identity $[a - c]_+ - [-a - c]_+ = a - c$ in general, which is why the result does not collapse to a linear law), gives the general damped dynamic of the main text, $\dot{s}_i = \gamma\kappa \sum_j s_i s_j \big([u_i - u_j - c]_+ - [u_j - u_i - c]_+\big)$. For two firms with $\Delta := u_1 - u_2 > 0$: the $j = 2$ term of $i = 1$'s sum has $u_2 - u_1 - c = -\Delta - c < 0$, so $[u_2 - u_1 - c]_+ = 0$, leaving $\dot{s}_1 = \gamma\kappa\, s_1 s_2\, [\Delta - c]_+$.

Taste-gap smoothing. Give each consumer an idiosyncratic match value ε per firm; a consumer currently at i meeting a sampled j-consumer faces the effective gap $(u_j - u_i) + G$ with $G = \varepsilon_j - \varepsilon_i$ drawn from a symmetric density g ($g(-v) = g(v)$), so the conditional switch probability is $\kappa[u_j - u_i + G - c]_+$ and the population flow requires averaging over the revising consumer's own draw of G. Define

$$h(\Delta) \;:=\; \mathbb{E}_G\big[[\Delta + G - c]_+\big] \;=\; \int_{c-\Delta}^{\infty} (\Delta + v - c)\, g(v)\, dv,$$

the integrand being zero below $v = c - \Delta$. The two-firm dynamic becomes $\dot{s}_1 = \gamma\kappa\, s_1 s_2\big(h(\Delta) - h(-\Delta)\big)$; setting g to a point mass at 0 recovers $h(\Delta) - h(-\Delta) = [\Delta - c]_+ - [-\Delta - c]_+ = [\Delta - c]_+$ exactly, the kinked law above, so the smoothed model nests it.

Differentiate h by Leibniz's rule with lower limit $L(\Delta) = c - \Delta$:

$$h'(\Delta) = -\underbrace{(\Delta + v - c)\big|_{v = L(\Delta)}}_{=\,\Delta + (c-\Delta) - c\,=\,0}\, g(L(\Delta))\, L'(\Delta) + \int_{L(\Delta)}^{\infty} 1 \cdot g(v)\, dv \;=\; \int_{c-\Delta}^{\infty} g(v)\, dv = \Pr[G \geq c - \Delta],$$

the boundary term vanishing identically because the integrand is zero exactly at the kink. At $\Delta = 0$: $h'(0) = \Pr[G \geq c] = \Pr[G > c]$ for continuous G. Because $h(\Delta) - h(-\Delta)$ is an odd function of Δ by construction, its Taylor series contains only odd powers, so $h(0) - h(-0) = 0$ exactly and the first correction past the linear term is $O(\Delta^3)$:

$$h(\Delta) - h(-\Delta) = 2\Delta\, h'(0) + O(\Delta^3) = 2\, \Pr[G > c]\, \Delta + O(\Delta^3).$$

The factor of two is the sum of two symmetric margins: the outflow from firm 1 and the inflow from firm 2 each contribute a sensitivity $\gamma\kappa \Pr[G > c]$ to a marginal change in Δ, and the reported statistic is the sensitivity of a single margin, $\beta_{\mathrm{D}}(c) := \gamma\kappa \Pr[G > c]$ — consistent with $\beta_{\mathrm{D}}(0) = \gamma\kappa \Pr[G > 0] = \gamma\kappa/2$, exactly half the frictionless $\beta_{\mathrm{D}} = \gamma\kappa$ of Proposition 2(a), since with symmetric taste noise and zero switching cost exactly half of all sampled pairs have an idiosyncratic gap that reinforces, rather than fights, the systematic ranking. To leading order in Δ,

$$\dot{s}_1 \;=\; \beta_{\mathrm{D}}(c)\, s_1 s_2\, \Delta, \qquad \beta_{\mathrm{D}}(c) = \gamma\kappa\, \Pr[G > c],$$

Proposition 3's smoothed law. Monotonicity is immediate from $g \geq 0$: $\Pr[G > c] = \int_c^\infty g(v)\, dv$ is non-increasing in c, so $\beta_{\mathrm{D}}(c)$ falls monotonically from $\gamma\kappa/2$ at $c = 0$ toward 0 as $c \to \infty$. □

B The discrete-time replicator as a Price equation

Section 3's Equation (2) asserts, without proof, that the discrete-time replicator is exactly the selection term of the Price equation with the transmission term set to zero. This appendix derives that claim, states the exact one-step identity underneath it, shows precisely where the $O(\Delta t^2)$ remainder in Equation (2) comes from, and fixes the notation that maps every quantity in this paper onto Paper No. 1's accounting [101].

The discrete map

Fix a time step Δt and define, for each firm i, a one-period fitness

$$w_i(t) = 1 + \Delta t\,\big(\pi_i(t) - \bar{\pi}(t)\big), \qquad \bar{\pi}(t) = \sum_j s_j(t)\, \pi_j(t),$$

and update shares by $s_i(t + \Delta t) = s_i(t)\, w_i(t)/\bar{w}(t)$, with $\bar{w}(t) = \sum_i s_i(t) w_i(t)$ the usual replicator-map normalization. Two facts follow immediately from the definition of w_i, before any approximation is made.

First, $\bar{w}(t) = 1$ exactly:

$$\bar{w}(t) = \sum_i s_i(t)\big[1{+}\Delta t(\pi_i(t){-}\bar{\pi}(t))\big] = \sum_i s_i(t){+}\Delta t\Big(\sum_i s_i(t)\pi_i(t){-}\bar{\pi}(t)\sum_i s_i(t)\Big) = 1{+}\Delta t\big(\bar{\pi}(t){-}\bar{\pi}(t)\big) = 1,$$

using $\sum_i s_i(t) = 1$ and $\sum_i s_i(t)\pi_i(t) = \bar{\pi}(t)$ by definition. A closed population — no entry or exit within the step — needs no renormalization at all: the division by $\bar{w}(t)$ that the general replicator map requires is here a division by one. Second, the update is therefore

$$s_i(t + \Delta t) = s_i(t)\, w_i(t) = s_i(t) + \Delta t\, s_i(t)\big(\pi_i(t) - \bar{\pi}(t)\big),$$

the forward-Euler discretization of the continuous replicator $\dot{s}_i = s_i(\pi_i - \bar{\pi})$ of Equation (1), evaluated at the coefficients $\pi_i(t), \bar{\pi}(t)$ prevailing at the start of the step.

The exact one-step Price identity

Paper No. 1's Identity 1 holds for any population with weights $s_i(t)$, fitness $w_i(t)$, and character $z_i(t)$ subject to a within-unit change Δz_i over the step [101]:

$$\Delta\bar{z} = \mathrm{Sel} + \frac{\mathbb{E}_s[w_i\, \Delta z_i]}{\bar{w}}, \qquad \mathrm{Sel} \equiv \frac{\mathrm{Cov}_s(w_i, z_i)}{\bar{w}}.$$

The replicator of Section 5 models pure reallocation: no firm's own trait z_i changes between revision rounds, only the shares attached to each fixed trait do, so $\Delta z_i = 0$ for every i — the gap Section 9 flags as the innovation break, left to the replicator-mutator extension of Paper No. 3. The transmission term therefore vanishes identically, and with $\bar{w} = 1$ from above,

$$\Delta\bar{z} = \mathrm{Cov}_s(w_i, z_i) = \mathrm{Cov}_s\left(1 + \Delta t(\pi_i - \bar{\pi}),\, z_i\right) = \Delta t\ \mathrm{Cov}_s(\pi_i, z_i),$$

the last step using that a covariance is unaffected by an additive constant and scales linearly with a multiplicative one, so the "1" and the firm-independent "$-\Delta t\,\bar{\pi}$" inside w_i both drop out. This is exact for one step at frozen coefficients: the discrete replicator literally is the Price equation's selection term with zero transmission, no approximation involved yet.

Where the $O(\Delta t^2)$ comes from

The frozen-coefficient step above is not the same object as the true continuous-time flow over an interval of length Δt, because in a real market π_i and $\bar{\pi}$ are functions of the share vector, which itself moves during $[t, t+\Delta t]$. Along the true flow, $\dot{\bar{z}} = \sum_i \dot{s}_i\, z_i = \sum_i s_i(\pi_i - \bar{\pi})z_i = \mathrm{Cov}_s(\pi_i, z_i)$ at every instant, so

$$\bar{z}(t + \Delta t) - \bar{z}(t) = \int_t^{t+\Delta t} \mathrm{Cov}_s[\tau]\big(\pi_i(\tau), z_i\big)\, d\tau.$$

A first-order Taylor expansion of the (differentiable) integrand around $\tau = t$ gives $\mathrm{Cov}_s[\tau](\pi_i(\tau), z_i) = \mathrm{Cov}_s[t](\pi_i(t), z_i) + O(\tau - t)$, and integrating this over an interval of width Δt turns the $O(\tau - t)$ remainder into $O(\Delta t^2)$: a term growing linearly from zero, integrated over a width-Δt window, sweeps out an area of order Δt^2. Hence

$$\Delta\mathbb{E}_s[z] = \mathrm{Cov}_s(\pi_i, z_i)\,\Delta t + O(\Delta t^2),$$

recovering Equation (2) exactly: the frozen-coefficient discrete replicator and Paper No. 1's exact per-step Price identity agree with the true continuous flow to first order in Δt, and disagree only at second order — the standard local truncation error of a forward-Euler step applied to a flow whose vector field is evaluated once, at the interval's start.

The map to Paper No. 1's notation

Three substitutions carry every result of this paper into Paper No. 1's accounting exactly. Paper No. 1's $w_i = s_i(t+1)/s_i(t)$, computed from realized shares with $\Delta t = 1$, is this paper's $w_i = 1 + \Delta t(\pi_i - \bar{\pi})$, specialized to a market in which profitability deviation is the only fitness-relevant quantity and left in continuous-Δt form so that $\Delta t \to 0$ recovers the ODE exactly. Paper No. 1's transmission term, generally positive wherever firms improve their own productivity between periods, is zero throughout Propositions 2 and 3: this paper's replicator only reallocates an existing distribution of traits, and never explains a firm becoming better at what it does, exactly the limitation Section 9 states as this analogy's third honest break. Paper No. 1's selection term $\mathrm{Sel} = \mathrm{Cov}_s(w_i, z_i)/\bar{w}$ has an analogous exact form for the combined channel: applying the identical argument above to $\dot{s}_i = \beta s_i(z_i - \bar{z})$ of Proposition 2(c) in place of $\dot{s}_i = s_i(\pi_i - \bar{\pi})$ gives

$$\mathrm{Sel} = \Delta t\,\beta\ \mathrm{Var}_s(z_i),$$

Regime	mean $\hat{\beta}$	s.d. of $\hat{\beta}_t$	first 10	last 10	HHI_T	top share_T
pure replicator ($\beta = 0.20$)	0.213	0.007	0.204	0.221	0.092	0.187
damped, $c = 0.20$	0.143	0.011	0.129	0.157	0.048	0.138
neutral drift, $\sigma = 0.10$	−0.004	0.030	−0.002	−0.009	0.040	0.125
replicator + noise	0.214	0.029	0.212	0.216	0.100	0.185

Table 3: Recovered selection coefficients across 40 periods. "First 10" and "last 10" are means over the first and last ten periods. The true selection intensity in regimes 1, 2 and 4 is $\beta = 0.20$; regime 2 damps it through the switching cost.

since $\mathrm{Cov}_s(z_i, z_i) = \mathrm{Var}_s(z_i)$: the market's measured selection intensity is the product of the combined channel intensity β and the population variance of the underlying efficiency trait — the market-competition form of Fisher's fundamental theorem that Metcalfe's synthesis first identified [49] and that Paper No. 1's empirical decomposition estimates directly from data without needing β and $\mathrm{Var}_s(z_i)$ separately [101].

C Simulation: what a selection regression recovers

This appendix reports a deterministic simulation (seed 20260830; script `sim_replicator.py` accompanies the source) whose only purpose is to show what the industrial-dynamics literature's standard "selection regression" recovers under each of the dynamics this paper distinguishes. One industry of $N = 100$ firms carries fixed log-efficiencies $z_i \sim \mathcal{N}(0, 0.30^2)$ and rank-size initial shares $s_i(0) \propto r_i^{-0.6}$ (initial HHI = 0.019), with size decoupled from efficiency by a random permutation. Four regimes run for $T = 40$ periods:

1. *pure replicator*: $w_i = 1 + \beta(z_i - \bar{z})$ with $\beta = 0.20$, where $\bar{z}$ is the share-weighted mean;
2. *damped replicator* (Proposition 3): pairwise flows $\Delta s_i = \beta \sum_j s_i s_j \big([z_i - z_j - c]_+ - [z_j - z_i - c]_+\big)$ with switching cost $c = 0.20$ — identical to regime 1 when $c = 0$;
3. *neutral drift*: $w_i = 1 + \varepsilon_i$ with Laplace-distributed ε_i of standard deviation $\sigma = 0.10$ (the heavy-tailed growth shocks of §7);
4. *replicator plus noise*: regimes 1 and 3 combined.

Each period the regression $g_i = a + \hat{\beta}\,(z_i - \bar{z}) + e_i$, with $g_i = \ln\big(s_i(t{+}1)/s_i(t)\big)$ and $\bar{z}$ the unweighted mean, is estimated by ordinary least squares — the specification of the productivity–growth regressions reviewed in §7.

Three readings follow. First, the regression recovers the pure replicator's β to within a few percent; the slight upward drift from 0.204 to 0.221 across the run is the second-order effect of log growth — $\partial \ln(1 + \beta x)/\partial x = \beta/(1 + \beta x)$ exceeds β for the majority of firms once the share-weighted mean $\bar{z}$ has moved above the unweighted mean, as it does under selection. Second, the switching cost of $c = 0.20$ (two-thirds of one standard deviation of z) attenuates the measured coefficient to 0.143, a factor of 0.67; the share of firm pairs whose efficiency gap exceeds c in this industry is 0.674. The two numbers coincide only approximately in general, but they carry the same message as Proposition 3: what the regression reports is $\beta_{\mathrm{D}}(c)$, the institutionally damped intensity, and it is smaller than the underlying β by roughly

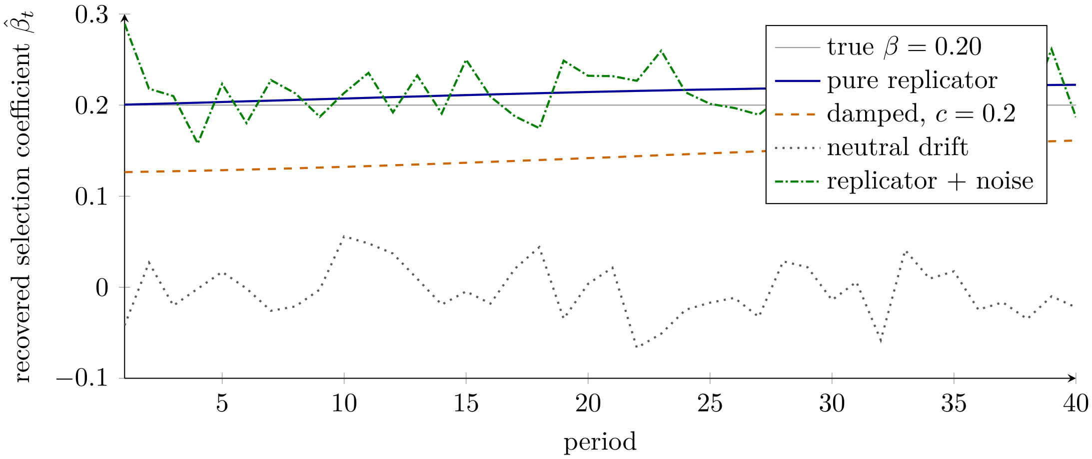

Figure 3: Period-by-period recovered $\hat{\beta}_t$ under the four regimes; the thin line marks the true $\beta = 0.20$. The pure replicator is recovered almost exactly; the switching cost removes roughly a third of the measured intensity; drift yields a coefficient indistinguishable from zero but with a standard error four times larger than in the noiseless regimes, which is the standard error a real panel would show.

the fraction of comparisons that fall inside the selection shadow. Third, adding realistic noise to the replicator (regime 4) leaves the mean coefficient intact but quadruples its dispersion: with $N = 100$ firms and Laplace-tailed shocks of $\sigma = 0.10$, a single cross-section carries a standard error of about 0.03 on a coefficient of 0.2, so the "weak selection" verdicts of §7 — coefficients that are positive, significant and small — are exactly what a damped replicator running inside heavy-tailed noise should produce.

The concentration columns record the other face of the same dynamics. Selection raises the Herfindahl index from 0.019 to 0.09–0.10 in forty periods and hands the best firm nearly a fifth of the market; the switching cost halves that concentration; pure drift doubles concentration with no efficiency gain whatsoever, the mechanism behind the drift null of the first paper in this series [101].

Figures 1 and 2 are generated by the same script: the two-strategy paths integrate $\dot{x} = x(1-x)\,[\Delta - c]_+$ and the smoothed law $\dot{x} = 2\,\Pr[G > c]\,x(1-x)\,\Delta$ with $\Delta = 0.30$, and the damping curves evaluate $2\,\Pr[G > c]$ for Gaussian and Laplace taste gaps as functions of c/τ, where τ is the standard deviation of a single consumer–firm match value.

References

[1] J. Maynard Smith, G. R. Price (1973). The Logic of Animal Conflict. *Nature, vol. 246, pp. 15–18*. https://public.ek-cer.hu/~szabo/EGT20/library/maynard_n73.pdf. DOI: 10.1038/246015a0.

[2] John Maynard Smith (1982). Evolution and the Theory of Games. *Cambridge University Press (book)*. https://archive.org/details/evolutiontheoryo0000mayn.

[3] Peter D. Taylor, Leo B. Jonker (1978). Evolutionarily Stable Strategies

and Game Dynamics. *Mathematical Biosciences, vol. 40, pp. 145–156.* https://static1.squarespace.com/static/5dc9917c4cdb1467e1e54c4c/t/6526a5647942e33b89ab7af1/1697031525653/1978+evolutionarily+stable+strategies+and+game+dynamics.pdf. DOI: 10.1016/0025-5564(78)90077-9.

[4] E. C. Zeeman (1980). Population Dynamics from Game Theory. *in "Global Theory of Dynamical Systems", Lecture Notes in Mathematics vol. 819 (Springer, Berlin/Heidelberg), pp. 471–497.* https://www.lms.ac.uk/sites/default/files/1980%20Population%20dynamics%20from%20game%20theory%20(preprint).pdf. DOI: 10.1007/BFb0087009.

[5] Peter Schuster, Karl Sigmund (1983). Replicator Dynamics. *Journal of Theoretical Biology, vol. 100, pp. 533–538.* https://doi.org/10.1016/0022-5193(83)90445-9. DOI: 10.1016/0022-5193(83)90445-9.

[6] Josef Hofbauer, Karl Sigmund (1998). Evolutionary Games and Population Dynamics. *Cambridge University Press (book).* https://archive.org/details/evolutionarygame0000hofb.

[7] Josef Hofbauer, Karl Sigmund (2003). Evolutionary Game Dynamics. *Bulletin of the American Mathematical Society, vol. 40, no. 4, pp. 479–519.* https://www.ams.org/journals/bull/2003-40-04/S0273-0979-03-00988-1/S0273-0979-03-00988-1.pdf. DOI: 10.1090/S0273-0979-03-00988-1.

[8] Jörgen W. Weibull (1995). Evolutionary Game Theory. *MIT Press (book).* https://archive.org/details/evolutionarygame0000weib.

[9] Larry Samuelson (1997). Evolutionary Games and Equilibrium Selection. *MIT Press, Economic Learning and Social Evolution series (book).* https://ideas.repec.org/b/mtp/titles/0262692198.html.

[10] William H. Sandholm (2010). Population Games and Evolutionary Dynamics. *MIT Press, Economic Learning and Social Evolution series (book).* https://openlibrary.org/books/OL26156681M/Population_Games_and_Evolutionary_Dynamics_Economic_Learning_and_Social_Evolution.

[11] Martin A. Nowak (2006). Evolutionary Dynamics: Exploring the Equations of Life. *Belknap Press of Harvard University Press (book).* https://openlibrary.org/books/OL7691827M/Evolutionary_Dynamics.

[12] Martin A. Nowak, Karl Sigmund (2004). Evolutionary Dynamics of Biological Games. *Science, vol. 303, pp. 793–799 (working-paper version opened: IIASA Interim Report IR-04-013, IIASA Studies in Adaptive Dynamics No. 85).* https://pure.iiasa.ac.at/id/eprint/7432/1/IR-04-013.pdf. DOI: 10.1126/science.1093411.

[13] Ross Cressman, Yi Tao (2014). The Replicator Equation and Other Game Dynamics. *Proceedings of the National Academy of Sciences (PNAS), vol. 111, Suppl. 3, pp. 10810–10817.* https://www.pnas.org/doi/10.1073/pnas.1400823111. DOI: 10.1073/pnas.1400823111.

[14] Karen M. Page, Martin A. Nowak (2002). Unifying Evolutionary Dynamics. *Journal of Theoretical Biology, vol. 219, no. 1, pp. 93–98*. https://doi.org/10.1006/jtbi.2002.3112. DOI: 10.1006/jtbi.2002.3112.

[15] Immanuel M. Bomze (1983). Lotka-Volterra Equation and Replicator Dynamics: A Two-Dimensional Classification. *Biological Cybernetics, vol. 48, pp. 201–211*. https://doi.org/10.1007/BF00318088. DOI: 10.1007/BF00318088.

[16] Arne Traulsen, Jens Christian Claussen, Christoph Hauert (2005). Coevolutionary Dynamics: From Finite to Infinite Populations. *Physical Review Letters, vol. 95, article 238701*. https://arxiv.org/abs/cond-mat/0409655. DOI: 10.1103/PhysRevLett.95.238701.

[17] Dean P. Foster, H. Peyton Young (1990). Stochastic Evolutionary Game Dynamics. *Theoretical Population Biology, vol. 38, pp. 219–232*. http://www.econ2.jhu.edu/people/Young/Stochastic_Evolutionary_Game_Dynamics.pdf.

[18] Michihiro Kandori, George J. Mailath, Rafael Rob (1993). Learning, Mutation, and Long Run Equilibria in Games. *Econometrica, vol. 61, no. 1, pp. 29–56*. http://www.dklevine.com/archive/refs4500.pdf.

[19] H. Peyton Young (1993). The Evolution of Conventions. *Econometrica, vol. 61, no. 1, pp. 57–84*. http://dklevine.com/archive/refs4485.pdf.

[20] Drew Fudenberg, David K. Levine (1998). The Theory of Learning in Games. *MIT Press, Economic Learning and Social Evolution series (book)*. https://books.google.com/books/about/The_Theory_of_Learning_in_Games.html?id=G6vTQFluxuEC.

[21] Karl H. Schlag (1998). Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-Armed Bandits. *Journal of Economic Theory, vol. 78, no. 1, pp. 130–156*. https://ideas.repec.org/p/els/esrcls/028.html. DOI: 10.1006/jeth.1997.2347.

[22] Jonas Björnerstedt, Jörgen W. Weibull (1994). Nash Equilibrium and Evolution by Imitation. *IUI (Industrial Institute for Economic and Social Research, Stockholm) Working Paper No. 407, January 1994; published as a chapter in "The Rational Foundations of Economic Behaviour" (K. J. Arrow et al., eds.), Macmillan/IUI, 1996, pp. 155–171*. https://www.ifn.se/wfiles/wp/wp407.pdf.

[23] Tilman Börgers, Rajiv Sarin (1997). Learning Through Reinforcement and Replicator Dynamics. *Journal of Economic Theory, vol. 77, no. 1, pp. 1–14*. http://www.dklevine.com/archive/refs4380.pdf. DOI: 10.1006/jeth.1997.2319.

[24] Ido Erev, Alvin E. Roth (1998). Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategy Equilibria. *American Economic Review, vol. 88, no. 4, pp. 848–881*. https://econweb.ucsd.edu/~jandreon/Econ264/papers/Erev%20Roth%20AER%201998.pdf.

[25] Lawrence E. Blume (1993). The Statistical Mechanics of Strategic Interaction. *Games and Economic Behavior, vol. 5, no. 3, pp. 387–424*. https://www.stat.berkeley.edu/~aldous/206-RWG/RWGpapers/blume.pdf. DOI: 10.1006/game.1993.1023.

[26] Itzhak Gilboa, Akihiko Matsui (1991). Social Stability and Equilibrium. *Econometrica, vol. 59, no. 3, pp. 859–867*. https://www.econometricsociety.org/publications/econometrica/1991/05/01/notes-and-comments-social-stability-and-equilibrium. DOI: 10.2307/2938230.

[27] George W. Brown (1951). Iterative Solution of Games by Fictitious Play. *in T. C. Koopmans (ed.), "Activity Analysis of Production and Allocation" (Cowles Commission Monograph 13), Chapman & Hall/Wiley, pp. 374–376*. https://archive.org/details/in.ernet.dli.2015.39951.

[28] Dov Monderer, Lloyd S. Shapley (1996). Potential Games. *Games and Economic Behavior, vol. 14, article no. 0044, pp. 124–143*. https://www.cs.tau.ac.il/~mansour/sem-game-02-03/monderer-potential-96.pdf. DOI: 10.1006/game.1996.0044.

[29] Michael J. Smith (1984). The Stability of a Dynamic Model of Traffic Assignment—An Application of a Method of Lyapunov. *Transportation Science, vol. 18, no. 3, pp. 245–252*. https://pubsonline.informs.org/doi/10.1287/trsc.18.3.245. DOI: 10.1287/trsc.18.3.245.

[30] George W. Brown, John von Neumann (1950). Solutions of Games by Differential Equations. *in H. W. Kuhn & A. W. Tucker (eds.), "Contributions to the Theory of Games I", Annals of Mathematics Studies no. 24, Princeton University Press, pp. 73–79; also circulated as RAND Corporation Paper P-142*. https://web.archive.org/web/20231208034320/https://www.rand.org/pubs/papers/P0142.html.

[31] Ed Hopkins (2002). Two Competing Models of How People Learn in Games. *Econometrica, vol. 70, no. 6, pp. 2141–2166*. https://ideas.repec.org/p/edn/esedps/51.html. DOI: 10.1111/j.1468-0262.2002.00436.x.

[32] Fernando Vega-Redondo (1997). The Evolution of Walrasian Behavior. *Econometrica, vol. 65, no. 2, pp. 375–384*. http://www.ivie.es/downloads/docs/wpasad/wpasad-1996-05.pdf. DOI: 10.2307/2171898.

[33] Mark E. Schaffer (1988). Evolutionarily Stable Strategies for a Finite Population and a Variable Contest Size. *Journal of Theoretical Biology, vol. 132, no. 4, pp. 469–478*. https://pubmed.ncbi.nlm.nih.gov/3226137/. DOI: 10.1016/S0022-5193(88)80085-7.

[34] Mark E. Schaffer (1989). Are Profit-Maximisers the Best Survivors? A Darwinian Model of Economic Natural Selection. *Journal of Economic Behavior & Organization, vol. 12, no. 1, pp. 29–45*. https://researchportal.hw.ac.uk/en/publications/are-profit-maximisers-the-best-survivors-a-darwinian-model-of-eco/. DOI: 10.1016/0167-2681(89)90075-9.

[35] Carlos Alós-Ferrer, Ana B. Ania (2005). The Evolutionary Stability of Perfectly Competitive Behavior. *Economic Theory, vol. 26, no. 3, pp. 497–516*. https://link.springer.com/article/10.1007/s00199-004-0474-8. DOI: 10.1007/s00199-004-0474-8.

[36] Jose Apesteguia, Steffen Huck, Jörg Oechssler (2007). Imitation—Theory and Experimental Evidence. *Journal of Economic Theory, vol. 136, no. 1, pp. 217–235*. https://jose-apesteguia.github.io/imitation.pdf. DOI: 10.1016/j.jet.2006.07.006.

[37] Steffen Huck, Hans-Theo Normann, Jörg Oechssler (1999). Learning in Cournot Oligopoly—An Experiment. *The Economic Journal, vol. 109, no. 454, pp. 80–95*. https://academic.oup.com/ej/article/109/454/80-95/5128699. DOI: 10.1111/1468-0297.00418.

[38] Theo Offerman, Jan Potters, Joep Sonnemans (2002). Imitation and Belief Learning in an Oligopoly Experiment. *Review of Economic Studies, vol. 69, no. 4, pp. 973–997*. https://academic.oup.com/restud/article-lookup/doi/10.1111/1467-937X.00233. DOI: 10.1111/1467-937X.00233.

[39] Daniel Friedman (1991). Evolutionary Games in Economics. *Econometrica, vol. 59, no. 3, pp. 637–666*. https://www.econometricsociety.org/publications/econometrica/1991/05/01/evolutionary-games-economics. DOI: 10.2307/2938222.

[40] George J. Mailath (1998). Do People Play Nash Equilibrium? Lessons from Evolutionary Game Theory. *Journal of Economic Literature, vol. 36, no. 3, pp. 1347–1374*. https://ideas.repec.org/a/aea/jeclit/v36y1998i3p1347-1374.html.

[41] Thorstein Veblen (1898). Why is Economics not an Evolutionary Science?. *The Quarterly Journal of Economics, Vol. 12, No. 4, pp. 373–397*. https://archive.org/details/jstor-1882952. DOI: 10.2307/1882952.

[42] Armen A. Alchian (1950). Uncertainty, Evolution, and Economic Theory. *Journal of Political Economy, Vol. 58, No. 3 (Jun. 1950), pp. 211–221*. https://ios23.classes.ryansafner.com/files/readings/Alchian-1950.pdf.

[43] Stephen Enke (1951). On Maximizing Profits: A Distinction Between Chamberlin and Robinson. *American Economic Review, Vol. 41, No. 4 (Sept. 1951), pp. 566–578*. https://archive.org/details/sim_american-economic-review_1951-09_41_4.

[44] Edith T. Penrose (1952). Biological Analogies in the Theory of the Firm. *American Economic Review, Vol. 42, No. 5 (Dec. 1952), pp. 804–819*. https://archive.org/details/sim_american-economic-review_1952-12_42_5.

[45] Milton Friedman (1953). The Methodology of Positive Economics. *Essays in Positive Economics, University of Chicago Press (book chapter)*. https://www.marxists.org/reference/subject/philosophy/works/us/friedman.htm.

[46] Sidney G. Winter (1964). Economic "Natural Selection" and the Theory of the Firm. *Yale Economic Essays, Vol. 4, pp. 225–272*. https://ideas.repec.org/h/ssa/lemchs/winter-1964.html.

[47] Richard R. Nelson, Sidney G. Winter (1982). An Evolutionary Theory of Economic Change. *Belknap Press of Harvard University Press (book), xi + 437 pp.*. https://archive.org/details/evolutionarytheo0000nels.

[48] Jack Downie (1958). The Competitive Process. *Duckworth (London), 1958, 199 pp., 25s. — confirmed via contemporaneous review by P. Lesley Cook, The Economic Journal, Vol. 69, No. 274 (Jun. 1959), pp. 360–361*. https://archive.org/details/sim_economic-journal_1959-06_69_274.

[49] J. Stanley Metcalfe (1994). Competition, Fisher's Principle and Increasing Returns in the Selection Process. *Journal of Evolutionary Economics, Vol. 4, No. 4, pp. 327–346*. https://ideas.repec.org/a/spr/joevec/v4y1994i4p327-46.html. DOI: 10.1007/BF01236409.

[50] J. Stanley Metcalfe (1998). Evolutionary Economics and Creative Destruction (The Graz Schumpeter Lectures). *Routledge, London (book)*. https://books.google.be/books?id=FwIBeI6EjLEC&redir_esc=y.

[51] Gerald Silverberg, Giovanni Dosi, Luigi Orsenigo (1988). Innovation, Diversity and Diffusion: A Self-Organisation Model. *The Economic Journal, Vol. 98, No. 393, pp. 1032–1054*. https://ideas.repec.org/a/ecj/econjl/v98y1988i393p1032-54.html. DOI: 10.2307/2233718.

[52] Giovanni Dosi, Richard R. Nelson (1994). An Introduction to Evolutionary Theories in Economics. *Journal of Evolutionary Economics, Vol. 4, No. 3, pp. 153–172*. https://ideas.repec.org/a/spr/joevec/v4y1994i3p153-72.html. DOI: 10.1007/BF01236366.

[53] Giovanni Dosi, Richard R. Nelson (2010). Technical Change and Industrial Dynamics as Evolutionary Processes. *Chapter 3, Handbook of the Economics of Innovation, Vol. 1 (B.H. Hall & N. Rosenberg, eds.), Elsevier, pp. 51–127 (working-paper version: LEM Working Paper 2009/07, Scuola Superiore Sant'Anna)*. https://www.econstor.eu/handle/10419/89555. DOI: 10.1016/S0169-7218(10)01003-8.

[54] Ulrich Witt (2008). What is Specific about Evolutionary Economics?. *Journal of Evolutionary Economics, Vol. 18, No. 5, pp. 547–575*. https://www.proquest.com/openview/fc3574a367fbc0841ebf0cf7cf4a524d/1?pq-origsite=gscholar&cbl=31401. DOI: 10.1007/s00191-008-0107-7.

[55] Geoffrey M. Hodgson, Thorbjørn Knudsen (2010). Darwin's Conjecture: The Search for General Principles of Social and Economic Evolution. *University of Chicago Press (book), 304 pp.*. https://press.uchicago.edu/ucp/books/book/chicago/D/bo10052998.html. DOI: 10.7208/chicago/9780226346922.001.0001.

[56] Howard E. Aldrich, Geoffrey M. Hodgson, David L. Hull, Thorbjørn Knudsen, Joel Mokyr, Viktor J. Vanberg (2008). In Defence of Generalized Darwinism. *Journal of Evolutionary Economics, Vol. 18, No. 5, pp. 577–596*. https://uhra.herts.ac.uk/id/eprint/80/. DOI: 10.1007/s00191-008-0110-z.

[57] Karolina Safarzynska, Jeroen C.J.M. van den Bergh (2010). Evolutionary Models in Economics: A Survey of Methods and Building Blocks. *Journal of Evolutionary Economics, Vol. 20, pp. 329–373*. https://research.vu.nl/en/publications/evolutionary-models-in-economics-a-survey-of-methods-and-building/. DOI: 10.1007/s00191-009-0153-9.

[58] Esben Sloth Andersen (2004). Population Thinking, Price's Equation and the Analysis of Economic Evolution. *Evolutionary and Institutional Economics Review, Vol. 1, No. 1, pp. 127–148*. https://vbn.aau.dk/en/publications/

population-thinking-prices-equation-and-the-analysis-of-economic-/. DOI: 10.14441/eier.1.127.

[59] Thorbjørn Knudsen (2002). Economic Selection Theory. *Journal of Evolutionary Economics, Vol. 12, No. 4, pp. 443–470*. https://ideas.repec.org/a/spr/joevec/v12y2002i4p443-470.html. DOI: 10.1007/s00191-002-0126-8.

[60] Geoffrey M. Hodgson (2019). Evolutionary Economics: Its Nature and Future. *Cambridge University Press, Elements in Evolutionary Economics series*. https://www.cambridge.org/core/elements/abs/evolutionary-economics/229EF6BA5434AB59ADCC1B3F706668BA. DOI: 10.1017/9781108767811.

[61] John Sutton (1997). Gibrat's Legacy. *Journal of Economic Literature, Vol. 35, No. 1, pp. 40–59*. https://econpapers.repec.org/RePEc:aea:jeclit:v:35:y:1997:i:1:p:40-59.

[62] M.H.R. Stanley, L.A.N. Amaral, S.V. Buldyrev, S. Havlin, H. Leschhorn, P. Maass, M.A. Salinger, H.E. Stanley (1996). Scaling Behaviour in the Growth of Companies. *Nature, Vol. 379, pp. 804–806*. https://www.nature.com/articles/379804a0. DOI: 10.1038/379804a0.

[63] Giulio Bottazzi, Angelo Secchi (2006). Explaining the Distribution of Firm Growth Rates. *The RAND Journal of Economics, Vol. 37, No. 2, pp. 235–256*. https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1756-2171.2006.tb00014.x. DOI: 10.1111/j.1756-2171.2006.tb00014.x.

[64] Giovanni Dosi, Daniele Moschella, Emanuele Pugliese, Federico Tamagni (2015). Productivity, Market Selection, and Corporate Growth: Comparative Evidence Across US and Europe. *Small Business Economics, Vol. 45, No. 3, pp. 643–672*. http://www.lem.sssup.it/WPLem/files/2013-15.pdf. DOI: 10.1007/s11187-015-9655-z.

[65] Giulio Bottazzi, Giovanni Dosi, Nadia Jacoby, Angelo Secchi, Federico Tamagni (2010). Corporate Performances and Market Selection: Some Comparative Evidence. *Industrial and Corporate Change, Vol. 19, No. 6, pp. 1953–1996*. http://www.lem.sssup.it/WPLem/files/2009-13.pdf. DOI: 10.1093/icc/dtq063.

[66] Uwe Cantner, Jens J. Krüger (2008). Micro-Heterogeneity and Aggregate Productivity Development in the German Manufacturing Sector — Results from a Decomposition Exercise. *Journal of Evolutionary Economics, Vol. 18, pp. 119–133*. https://link.springer.com/article/10.1007/s00191-007-0079-z. DOI: 10.1007/s00191-007-0079-z.

[67] Lucia Foster, John Haltiwanger, Chad Syverson (2008). Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?. *American Economic Review, Vol. 98, No. 1, pp. 394–425 (circulated as NBER Working Paper 11555)*. https://www.nber.org/papers/w11555. DOI: 10.1257/aer.98.1.394.

[68] Lucia Foster, John Haltiwanger, C.J. Krizan (2001). Aggregate Productivity Growth: Lessons from Microeconomic Evidence. *Chapter 8 in "New Developments in Productivity Analysis" (C.R. Hulten, E.R. Dean, M.J. Harper, eds.), University of Chicago Press*

for NBER (circulated as NBER Working Paper 6803, 1998). https://www.nber.org/papers/w6803.

[69] G. Steven Olley, Ariel Pakes (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica, Vol. 64, No. 6, pp. 1263–1297 (circulated as NBER Working Paper 3977).* https://www.nber.org/papers/w3977. DOI: 10.2307/2171831.

[70] Marc J. Melitz, Sašo Polanec (2015). Dynamic Olley-Pakes Productivity Decomposition with Entry and Exit. *RAND Journal of Economics, Vol. 46, No. 2, pp. 362–375 (circulated as NBER Working Paper 18182).* https://www.nber.org/papers/w18182. DOI: 10.1111/1756-2171.12088.

[71] Eric J. Bartelsman, John C. Haltiwanger, Stefano Scarpetta (2013). Cross-Country Differences in Productivity: The Role of Allocation and Selection. *American Economic Review, Vol. 103, No. 1, pp. 305–334 (circulated as NBER Working Paper 15490).* https://www.nber.org/papers/w15490. DOI: 10.1257/aer.103.1.305.

[72] Chad Syverson (2004). Product Substitutability and Productivity Dispersion. *The Review of Economics and Statistics, Vol. 86, No. 2, pp. 534–550 (circulated as NBER Working Paper 10049, Oct. 2003).* https://www.nber.org/papers/w10049. DOI: 10.1162/003465304323031094.

[73] Chad Syverson (2011). What Determines Productivity?. *Journal of Economic Literature, Vol. 49, No. 2, pp. 326–365.* https://www.aeaweb.org/articles?id=10.1257/jel.49.2.326. DOI: 10.1257/jel.49.2.326.

[74] Ryan Decker, John Haltiwanger, Ron Jarmin, Javier Miranda (2014). The Role of Entrepreneurship in US Job Creation and Economic Dynamism. *Journal of Economic Perspectives, Vol. 28, No. 3, pp. 3–24.* https://econweb.umd.edu/~haltiwan/jep_dhjm.pdf. DOI: 10.1257/jep.28.3.3.

[75] Ryan A. Decker, John Haltiwanger, Ron S. Jarmin, Javier Miranda (2020). Changing Business Dynamism and Productivity: Shocks Versus Responsiveness. *American Economic Review, Vol. 110, No. 12, pp. 3952–3990 (circulated as NBER Working Paper 24236, Jan. 2018).* https://www.nber.org/papers/w24236. DOI: 10.1257/aer.20190680.

[76] Chang-Tai Hsieh, Peter J. Klenow (2009). Misallocation and Manufacturing TFP in China and India. *Quarterly Journal of Economics, Vol. 124, No. 4, pp. 1403–1448 (circulated as NBER Working Paper 13290).* https://www.nber.org/papers/w13290. DOI: 10.1162/qjec.2009.124.4.1403.

[77] Zvi Griliches (1957). Hybrid Corn: An Exploration in the Economics of Technological Change. *Econometrica, Vol. 25, No. 4, pp. 501–522.* https://www.jstor.org/stable/1905380. DOI: 10.2307/1905380.

[78] Frank M. Bass (1969). A New Product Growth for Model Consumer Durables (commonly cited as “A New Product Growth Model for Consumer Durables”). *Management Science, Vol. 15, No. 5, pp. 215–227.* https://pubsonline.informs.org/doi/10.1287/mnsc.15.5.215. DOI: 10.1287/mnsc.15.5.215.

[79] H. Peyton Young (2009). Innovation Diffusion in Heterogeneous Populations: Contagion, Social Influence, and Social Learning. *American Economic Review, Vol. 99, No. 5, pp. 1899–1924*. https://www.aeaweb.org/articles?id=10.1257/aer.99.5.1899. DOI: 10.1257/aer.99.5.1899.

[80] Paul A. Geroski (2000). Models of Technology Diffusion. *Research Policy, Vol. 29, No. 4–5, pp. 603–625*. https://ideas.repec.org/a/eee/respol/v29y2000i4-5p603-625.html. DOI: 10.1016/s0048-7333(99)00092-x.

[81] Paul Klemperer (1987). Markets with Consumer Switching Costs. *The Quarterly Journal of Economics, Vol. 102, No. 2, pp. 375–394*. https://www.czaj.org/pub/teaching/IO/Markets%20with%20Consumer%20Switching%20Costs.pdf. DOI: 10.2307/1885068.

[82] Paul Klemperer (1995). Competition when Consumers have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade. *The Review of Economic Studies, Vol. 62, No. 4, pp. 515–539*. https://www.nuffield.ox.ac.uk/users/klemperer/competition.pdf. DOI: 10.2307/2298075.

[83] Joseph Farrell, Paul Klemperer (2007). Coordination and Lock-In: Competition with Switching Costs and Network Effects. *Handbook of Industrial Organization, Vol. 3, ch. 31, pp. 1967–2072 (Elsevier)*. http://www.nuff.ox.ac.uk/users/klemperer/Farrell_KlempererWP.pdf. DOI: 10.1016/s1573-448x(06)03031-7.

[84] Michael L. Katz, Carl Shapiro (1985). Network Externalities, Competition, and Compatibility. *American Economic Review, Vol. 75, No. 3, pp. 424–440*. https://idv.sinica.edu.tw/kongpin/teaching/io/KatzShapiro1.pdf.

[85] W. Brian Arthur (1989). Competing Technologies, Increasing Returns, and Lock-In by Historical Events. *The Economic Journal, Vol. 99, No. 394, pp. 116–131*. http://rochelleterman.com/ir/sites/default/files/arthur%201989.pdf. DOI: 10.2307/2234208.

[86] Carl Shapiro, Hal R. Varian (1998). Information Rules: A Strategic Guide to the Network Economy. *Harvard Business School Press (book)*. https://www.ischool.berkeley.edu/research/publications/1998/information-rules-strategic-guide-network-economy.

[87] William J. Baumol, John C. Panzar, Robert D. Willig (1982). Contestable Markets and the Theory of Industry Structure. *Harcourt Brace Jovanovich (book)*. https://books.google.com/books/about/Contestable_Markets_and_the_Theory_of_In.html?id=hHdgAAAAIAAJ.

[88] Nicholas Bloom, John Van Reenen (2007). Measuring and Explaining Management Practices Across Firms and Countries. *The Quarterly Journal of Economics, Vol. 122, No. 4, pp. 1351–1408*. https://worldmanagementsurvey.org/wp-content/images/2010/07/Measuring-and-Explaining-Management-Practices-Across-Firms-and-Countries-Bloom-and-Van-pdf. DOI: 10.1162/qjec.2007.122.4.1351.

[89] Nicholas Bloom, Benn Eifert, Aprajit Mahajan, David McKenzie, John Roberts (2013). Does Management Matter? Evidence from India. *The Quarterly Journal of Economics, Vol. 128, No. 1, pp. 1–51*. https://poverty-action.org/sites/default/files/publications/Does-Management-Matter-Evidence-from-India-Bloom-et-al-2013.pdf. DOI: 10.1093/qje/qjs044.

[90] Michela Giorcelli (2019). The Long-Term Effects of Management and Technology Transfers. *American Economic Review, Vol. 109, No. 1, pp. 121–152*. https://www.aeaweb.org/articles?id=10.1257/aer.20170619. DOI: 10.1257/aer.20170619.

[91] Everett M. Rogers (2003). Diffusion of Innovations, 5th edition. *Free Press, New York (book)*. https://openlibrary.org/books/OL15577453M/Diffusion_of_innovations.

[92] Robert L. Axtell (2001). U.S. Firm Sizes are Zipf Distributed (author's preprint circulated under this title; published as "Zipf Distribution of U.S. Firm Sizes"). *Science, Vol. 293, No. 5536, pp. 1818–1820 (author's Feb. 22, 2001 preprint, originally submitted to Nature, opened here; published in Science Sept. 2001)*. https://faculty.sites.iastate.edu/tesfatsi/archive/tesfatsi/USFirmSizesAreZipfDistributed.RAxtell2001.pdf. DOI: 10.1126/science.1062081.

[93] Xavier Gabaix (2009). Power Laws in Economics and Finance. *Annual Review of Economics, Vol. 1, pp. 255–294 (circulated as NBER Working Paper No. 14299)*. https://www.nber.org/system/files/working_papers/w14299/w14299.pdf. DOI: 10.1146/annurev.economics.050708.142940.

[94] Erzo G.J. Luttmer (2007). Selection, Growth, and the Size Distribution of Firms. *The Quarterly Journal of Economics, Vol. 122, No. 3, pp. 1103–1144*. http://users.econ.umn.edu/~luttmer/research/fd/Luttmer_QJE.pdf. DOI: 10.1162/qjec.122.3.1103.

[95] Hugo A. Hopenhayn (1992). Entry, Exit, and Firm Dynamics in Long Run Equilibrium. *Econometrica, Vol. 60, No. 5, pp. 1127–1150*. https://www.econometricsociety.org/publications/econometrica/1992/09/01/entry-exit-and-firm-dynamics-long-run-equilibrium. DOI: 10.2307/2951541.

[96] Boyan Jovanovic (1982). Selection and the Evolution of Industry. *Econometrica, Vol. 50, No. 3, pp. 649–670*. https://www.econometricsociety.org/publications/econometrica/1982/05/01/selection-and-evolution-industry. DOI: 10.2307/1912606.

[97] Richard Ericson, Ariel Pakes (1995). Markov-Perfect Industry Dynamics: A Framework for Empirical Work. *Review of Economic Studies, Vol. 62, No. 1, pp. 53–82*. http://aguirregabiria.net/courses/eco2901/ericson_pakes.pdf. DOI: 10.2307/2297841.

[98] Martin A. Nowak, Akira Sasaki, Christine Taylor, Drew Fudenberg (2004). Emergence of Cooperation and Evolutionary Stability in Finite Populations. *Nature, Vol. 428, pp. 646–650*. https://www.nature.com/articles/nature02414. DOI: 10.1038/nature02414.

[99] Dirk Helbing (1992). A Mathematical Model for Behavioral Changes by Pair Interactions. *in "Economic Evolution and Demographic Change" (Lecture Notes in Economics*

and Mathematical Systems, vol. 395), Springer Berlin Heidelberg, pp. 330–348. https://link.springer.com/chapter/10.1007/978-3-642-48808-5_18. DOI: 10.1007/978-3-642-48808-5_18.

[100] Alan Kirman (1993). Ants, Rationality, and Recruitment. *The Quarterly Journal of Economics, Vol. 108, No. 1, pp. 137–156.* https://fbaum.unc.edu/teaching/PLSC541_Fall06/Kirman%201993%20QJE.pdf. DOI: 10.2307/2118498.

[101] Brecht Corbeel (2026). Selection Accounting: The Price Equation as the Fundamental Theorem of Evolutionary Economics. *The Evolutionary Economics Papers, No. 1, Absolute Digital Publishers.* https://absolutedigitalpublishers.com/papers/evolutionary-economics-01-selection-accounting.pdf.