Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Micro-Heuristics, Macro-Strategy: An Agent-Based Model of Contingent Build, Partner, and Acquire Choice in Legacy-to-Emerging Market Transitions

Yingzheng Liu

yliu2@trinity.edu

University of Virginia

Shun Cao

Binghamton University

Zhen Liu

University of Virginia

Suning Zhu

Trinity University

Research Article

Keywords: capability sourcing, industry evolution, evolutionary economics, agent-based modeling, dynamic capabilities, first-mover advantage

Posted Date: October 30th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11166180/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Micro-Heuristics, Macro-Strategy:
An Agent-Based Model of Contingent Build, Partner, and Acquire Choice in Legacy-to-Emerging Market Transitions

Yingzheng Liu[1,2,*], Shun Cao[3], Zhen Liu[2], and Suning Zhu[1]

[1]Department of Finance and Business Analytics, Trinity University, San Antonio, TX, USA
[2]Department of Civil and Environmental Engineering, University of Virginia, Charlottesville, VA, USA
[3]School of Management, Binghamton University, Binghamton, NY, USA

Emails: yliu2@trinity.edu; scao12@binghamton.edu; rtv7bq@virginia.edu; szhu1@trinity.edu
[*]Corresponding author: Yingzheng Liu, yliu2@trinity.edu

Abstract

An established firm whose core market is declining while a new one opens must decide how to source the capabilities it demands: build them internally, partner with a firm that has them, acquire such a firm, or wait. We study this decision with an agent-based model of industry evolution in which heterogeneous, boundedly rational incumbents follow evidence-based routines and face market selection while crossing from a declining legacy market into a growing one. Because the model describes a generic transition, we use it for *generative foresight*: reasoning forward from micro-rules to the market structures they make likely, and testing its headlines by removing a supposed cause, randomizing an alternative cause, or decomposing its channel. Rivals' acquisitions change which routes are feasible, and the route mix follows feasibility rather than preference. In the plane spanned by the costs of building and partnering, build, partner, and acquire each produce the most market leaders in a region; per entrant, however, acquisition at its baseline cost leads in most of that plane, partnering keeps the region where it is cheap, and build's region reflects how often firms build. Early entry carries a time-in-market advantage that survives randomizing entry order and stays positive at every pace tested. Among entrants, acquisition has the highest end-of-run leader-or-survivor rate at the calibrated build scale; partnering overtakes it when builds start larger. The model separates three questions that sourcing comparisons often merge: which routes are feasible, which produces the most leaders, and how parents entering by each route fare.

Keywords: capability sourcing; industry evolution; evolutionary economics; agent-based modeling; dynamic capabilities; first-mover advantage

JEL classification: B52 (Institutional; Evolutionary); C63 (Computational Techniques; Simulation Modeling); D83 (Search; Learning; Information and Knowledge); L10 (Market Structure, Firm Strategy, and Market Performance); O33 (Technological Change: Choices and Consequences; Diffusion Processes)

1 Introduction

Firms in many industries now face disruption that does not let up. Advances in artificial intelligence and digital platforms, shifts in what customers want, and geopolitical realignment can all, at the same time, erode an old market and open a new one (Vial, 2019; Hanelt et al., 2021; Warner and Wäger, 2019; Brewis et al., 2023). Film giving way to digital imaging, the internal combustion engine to the electric vehicle, and on-premise software to cloud computing are three such transitions, and we return to each of them when we read the model's results against the record. Because these forces arrive together, an incumbent rarely gets a clean choice between a stable old market and a well-defined new one. It has to cross from the old world into the new while both are still changing, and it has to keep enough of the legacy business alive to survive the crossing (Verhoef et al., 2021; Eggers and Park, 2018).

Firms hit by the same disruption do not share the same fate. Some reposition and go on to lead the new market; others fail even though they saw the threat coming early (Eggers and Park, 2018; Tripsas and Gavetti, 2000). That gap in outcomes is what our study is about, and the central problem it poses is how incumbents respond.

How to respond is not a single choice. It is a multi-step process in which early moves shape later ones, and dynamic capability theory gives us a natural way to organize it (Teece et al., 1997; Teece, 2007). In that theory, renewal runs through three linked activities. A firm first *senses* that the legacy business is eroding and a new opportunity is opening. It then *seizes* the opportunity by committing to a route: internal development, a partnership, an acquisition, or a deliberate wait. Finally, it *transforms* itself to absorb and scale the new capability. These activities are not separate. What each one is worth depends on how the firm links them over time, and that interdependence is where the difficulty lies. Dynamic capabilities organize the firm's decision stages; the evolutionary perspective explains how repeated decisions, competition, and exit change the population of firms (Nelson and Winter, 1982). We use the first to specify what each firm does and the second to follow what the population becomes.

Prior research, though, gives a fragmented account of this process: there is a gap in each of the three stages and a fourth in how the stages fit together. The first gap is in sensing. Studies of environmental scanning explain when firms notice change (Cyert and March, 1963; Greve, 2003), but say little about the next step: when noticing a change turns into actually spending attention and money on one route. We know little about when a firm moves from seeing the disruption to choosing among building, partnering, acquiring, and waiting.

The second gap is in seizing. The build–borrow–buy framework, in which a firm can build a capability internally, borrow it through a partnership, or buy it through an acquisition, explains how these routes trade off speed, control, and access to knowledge (Capron and Mitchell, 2009, 2012). Innovation-management research reaches the same three routes from the R&D side, as technology sourcing: the choice between internal R&D and external knowledge acquisition, which turns out to be a choice of mix rather than of one over the other (Pisano, 1990; Veugelers, 1997; Cassiman and Veugelers, 2006; West and Bogers, 2014). The closest empirical work treats the choice among acquisitions, alliances, and divestitures as a one-time snapshot (Villalonga and McGahan, 2005). A separate body of work on entry timing explains when moving early or late pays (Lieberman and Montgomery, 1988; Suarez and Lanzolla, 2007; Zachary et al., 2015). The two are rarely joined, even though a route is worth more or less depending on when the firm takes it. Studying route choice and timing apart therefore leaves open when a firm should commit and how that timing shapes where it ends up.

The third gap is in transforming. Whether a firm can actually pull off a route depends on the resources it starts with and on how well it can rearrange its assets after entry (Cohen and Levinthal, 1990; Lane and Lubatkin, 1998). Prior work pays little attention to how these differences decide which routes are open to a given firm, and how well each one then goes. A firm that ignores these limits can commit to a route it cannot execute.

The fourth gap is in how the three stages fit together, because disruption response feeds back on itself in two directions. Inside a firm, sensing shapes seizing, seizing shapes transformation, and transformation changes what the firm senses next. Across firms, one firm's move changes the next firm's options: an acquisition takes a target off the market, an alliance can help set a standard, and an early internal build raises the bar later entrants must clear (Arthur, 1989; Klepper, 1996). The value of each route therefore depends on the choices the whole population is making and on how fast the technology frontier is moving. Existing work rarely asks how all this adds up to market-level outcomes: which route ends up dominant, when the disruption locks in, and how concentrated the new market becomes.

Together the gaps lead to one question.

Central question. *During a legacy-to-emerging market transition, how do market conditions and rivals' choices shape whether incumbents build, partner, acquire, or wait, and how do these routes compare in the leaders they produce and in the end-of-run outcomes of the parents that enter by each?*

We break it into four subquestions.

SQ1: *Sensing.* Do signals of legacy erosion and emerging opportunity change which route a firm takes by changing what it *wants*, or by changing what rivals' earlier choices leave *feasible*?

SQ2: *Timing.* How does the timing of a commitment, relative to the disruption, shape long-run leadership, and how does that advantage change with the pace of the transition?

SQ3: *Resources.* How do firm-level resources sort incumbents across the routes, and is the resulting spread in outcomes inherited from where firms start or produced as the system runs?

SQ4: *Aggregation.* How do these individual choices add up to a market structure (the route mix, the timing of consolidation, and the level of concentration), and which basic features set the boundaries between the mode-dominant regions?

To answer these questions, we build an agent-based model (ABM). In an ABM, many different firms each follow their own decision rules and interact, and the large-scale patterns grow out of those interactions instead of being built in (Bonabeau, 2002; Harrison et al., 2007; Axtell and Farmer, 2025). Management research has used the method to revisit March's exploration and exploitation model, to model strategic issue diagnosis by top management teams, and to study team composition and routine dynamics (Chanda and Miller, 2019; Miller and Lin, 2022; Blanco-Fernández et al., 2025; Gao and Akbaritabar, 2022). An agent-based model is what the fourth gap needs. The gap is about feedback, between stages and between firms, and the model represents that feedback directly: an acquisition removes a target from the pool every later decision-maker draws on, an affiliate raises the barrier later entrants face, and each year's outcomes set the conditions the next year's decisions read. Our model places incumbents in two linked markets, a declining legacy market and a growing emerging market. Each year every incumbent reads the market conditions, and a

firm whose attention is caught then chooses whether to build, partner, acquire, or wait. Year by year, these individual choices add up to a market structure.

Because the model describes a generic transition and not one industry, we do not fit it to a historical panel. We use it generatively instead, in the tradition of generative social science and of history-friendly modeling (Epstein, 1999; Davis et al., 2007; Malerba et al., 1999): we reason forward from evidence-based rules about how firms behave to the market structures those rules tend to produce, with every assumption out in the open. We call this use *generative foresight*. It shares with exploratory modeling under deep uncertainty the aim of reasoning about possible futures where projecting from past data breaks down (Bankes, 1993; Kwakkel and Pruyt, 2013), and research on corporate foresight treats that kind of reasoning as one of the building blocks of dynamic capabilities (Rohrbeck and Kum, 2018; Schwarz et al., 2020).

Two devices discipline the exercise. First, the model is still held to evidence, just evidence of a different kind: its rules follow qualitative patterns and parameter ranges taken from prior studies, and two of its core settings are tuned so that the baseline reproduces well-documented facts about how such markets behave. Second, each generative, confound, and decomposition headline is paired with a manipulation of its mechanism, a *mechanism knockout*: we switch off, randomize, or decompose the mechanism that is supposed to produce a pattern and check whether the pattern behaves as its type requires; the leader-or-survivor ranking is tested instead by sweeping the scale at which builds start. What this discipline produces is a *phase map*: a chart of which sourcing route produces the most leaders under which combination of conditions.

The central finding is a contingency in route choice, qualified by what happens to entrants. In the plane spanned by the costs of building and partnering, build, partner, and acquire each produce the most leaders in a region, and, in a separate sweep, acquisition's region gives way to partnering as the cost of acquisition rises above its baseline. Waiting is a choice every year; it enters the analysis through the share of incumbents that never enter and through the feasibility and preference measures, but it produces no leaders and so wins no region. Scored per entrant rather than by the number of leaders, acquisition leads over most of the plane, partnering keeps the region where it is cheap, and build's region reflects how often firms build. Two further findings qualify the picture. A first-mover advantage operates as a time-in-market effect that survives randomizing entry order; in the timing experiment, with the imitation and obsolescence channels switched on, it narrows at the fastest paces when firms time their own entry, but not when timing or order is randomized. And at the calibrated build scale, parents whose first entry was through acquisition most often end as leaders or survivors; parents entering first through partnership overtake them when internal builds start at a larger scale.

The study makes three contributions. First, **a dynamic theory of capability sourcing.** We take the build–partner–acquire contingency that resource-pathways theory states statically and let it move: the menu of routes is drained by rivals' choices, a route's value depends on its timing against a moving frontier, and the route a firm can best *take* comes apart from the route that best *survives*. Second, **a micro-to-macro account of market transition.** We show how many firms' individual sourcing rules add up to a market structure, a map of which route produces the most leaders where, a shakeout, and rising concentration, with concentration and the shakeout each collapsing when the rule that generates it is switched off. Third, **a generative-foresight method.** Each generative, confound, and decomposition headline is held to the result its type requires (collapse or a sharp fall when its rule is removed, persistence when an alternative cause is randomized, a channel split under decomposition); this extends the switch-off experiments that are standard in evolutionary simulation (Malerba et al., 1999; Dawid, 2006) to a model with no

historical data to fit. Section 5 draws out what these comparisons imply for managers.

The remainder of the paper proceeds as follows. Section 2 develops the theoretical framework and derives four refutable propositions. Section 3 specifies the model in full in the body of the paper, element by element, pairs each element with the phenomenon it stands for, justifies the design choices, and describes the experiments. Section 4 reports the results and examines them through the corresponding manipulations and parameter sweeps. Section 5 discusses the theoretical, methodological, and managerial contributions, Section 6 states the limitations, and Section 7 concludes.

2 Dynamic capabilities as the theoretical framework

Dynamic capability theory explains how firms adapt to fast-changing environments. It stresses a firm's ability to sense opportunities and threats, to seize opportunities through committed action, and to transform its resource base over time (Teece et al., 1997; Teece, 2007; Pitelis et al., 2024). Under steady disruption, these capabilities are not abstract traits but everyday processes, because they decide how a firm spots change, commits to a route, and rebuilds its operations (Schilke et al., 2018; Helfat and Raubitschek, 2018). Building on this view, we treat sensing, seizing, and transforming as interdependent processes: they are carried out through sequences of decisions, they change over time, and they differ from firm to firm. This process view is what lets us study how different firms respond differently to the same disruption, and how their choices together shape the market.

We take the three capabilities in turn. For each, we first review the theory, then explain how the model represents it, and finally state the testable expectation that follows. This three-step structure also previews how the model in Section 3 is organized, so that no part of the model arrives before the theory that motivates it.

2.1 Sensing: recognizing when the legacy core is no longer enough

Sensing is the ability to spot signals of change, including the decline of an existing market and the rise of a new one. But a signal being out there does not mean every firm picks it up. Firms search their environment with different intensity, and what drives that difference is performance feedback. Research shows that search intensifies when performance stays below the firm's aspirations, when market share slips, and when familiar routines no longer explain what is happening (Cyert and March, 1963; Greve, 2003).

Even when a firm detects a signal, detecting it is not the same as acting on it. Turning a signal into real strategic attention runs into both structural and cognitive obstacles. Dependence on existing customers, entrenched routines, prior commitments, and plain organizational inertia all tilt managers' attention toward the old market and delay their recognition of the new one (Christensen and Bower, 1996; Gilbert, 2005; Tripsas and Gavetti, 2000). Sensing is therefore selective, and it depends on the firm's history. Bad news pulls attention toward the new market, and inertia holds it on the legacy core.

In this study, we model sensing as an attention process that acts as a gate on entry. A firm's attention to the emerging market rises with its performance gap, its loss of legacy share, the distance to the new technological frontier, a disruption signal, and its own prior entry; it falls with dependence on existing customers, age, and low adaptability. Attention decides only whether a

firm considers entry at all, not which route it then takes. We keep the two separate. If attention also carried a preference for one route, we could no longer tell who searches from what searchers choose, and the first subquestion depends on keeping those two apart.

This separation gives us the first testable expectation. If sensing governs only who searches and when, then the mix of routes firms actually take should shift with conditions that change what is within reach, not with a fixed taste for one route. The test therefore separates a route's *feasibility*, whether the route is within reach for the firm, from its *preference*, whether the firm wants it. Both measures are defined in Section 3.6.

2.2 Seizing: build, partner, or acquire as alternative commitment logics

Once a firm sees the opportunity, it must seize it by committing to a route. We use the term seizing in a specific, working sense: the choice of which route to commit to. This is narrower than Teece's emphasis on investment and business-model design, and closer to the governance-mode choice studied in the resource-pathways and make-or-ally literatures. We use it this way because the route decision is what our central question is about, while acknowledging that the boundaries of the dynamic-capabilities construct are themselves debated (Eisenhardt and Martin, 2000; Helfat et al., 2007; Schilke and Helfat, 2025).

The build–borrow–buy framework holds that the three routes offer different mixes of speed, control, flexibility, and access to knowledge (Capron and Mitchell, 2009, 2012). Internal development keeps control and builds deep capability, but it starts slowly and is hard to coordinate. Partnership stays flexible and keeps options open under uncertainty, but it is worth only as much as the partner fits, agrees to the deal, and can be trusted not to leak too much knowledge (Folta, 1998; Kogut, 1991; Dyer and Singh, 1998). Acquisition reaches the market fast and locks up scarce assets, but its payoff depends on the price paid, the relative size of the two firms, and how costly it is to integrate a distant knowledge base (Haspeslagh and Jemison, 1991; Ahuja and Katila, 2001; Puranam et al., 2006). The technology-sourcing literature reaches the same three routes from the R&D side, as the choice among internal R&D, external technology acquisition, and collaborative development, and finds them to be complements as often as substitutes (Veugelers, 1997; Cassiman and Veugelers, 2006).

But seizing is more than picking a route from a menu, because the right route depends on the firm's resources and on the changing market. The same opportunity can lead one firm to build, another to partner, a third to acquire, and a fourth to wait, because firms differ in their resources and in how well each route fits the conditions they face. This is the theoretical root of our contingency thesis. No route is best for every firm, because a route's value comes from how the firm's position meets the market.

Put this way, the contingency is not itself new. The resource-pathways framework is already a contingency theory: its decision logic is a rule for picking the route that fits a firm's conditions, based on how relevant the firm's internal resources are, how tradable the target capability is, how close its partner relations are, and how easily it can integrate (Capron and Mitchell, 2012). Our contribution is therefore not the discovery that route choice is contingent. It is to let that contingency emerge from the firms' interactions and change over time, in three ways the static framework does not reach. First, the menu of choices is itself in play: a rival's acquisition removes a target from the pool, so the contingency works through what is still available, not only through what a manager would prefer. Second, a route's value depends on its timing against a moving frontier, which ties the route decision to the entry-timing problem the framework treats separately.

Third, the route a firm can best take and the route that best survives need not be the same, a split that a framework focused on the one-time choice does not address. The model is thus a dynamic, interaction-aware extension of the resource-pathways logic.

Seizing is also about when to commit, not only which route to take, and a route's value depends on its timing relative to the disruption. Committing early can win a first-mover advantage but faces high uncertainty; committing later faces less uncertainty but more competition. Product life cycle simulations show that whether entering pays depends on whether the entrant aims at quick profit or at long-run survival (Werker, 2003). A firm therefore weighs which route to take together with when to take it.

In this study, we model seizing as a scored choice. We compute each route's value from one shared scoring formula, subtract a cost specific to that route, and let the firm choose by a softmax rule, a probabilistic rule under which a route is picked more often the higher its score, without assuming the firm optimizes perfectly. The valuation formula and the outcome thresholds are the same for every route; the routes differ in their costs, in the starting positions they deliver, and in their post-entry transfer, dividend, and patience rules, each stated with its value in Section 3. This design gives two further expectations. First, because firms differ in resources, commitment should sort by *slack*, a firm's cash cushion measured against a fixed scale: resource-rich firms should lean toward the control-heavy routes and resource-poor firms toward the flexible ones. Second, because one firm's commitment changes what is left for the next, the most common route at any moment should depend on the choices already made.

2.3 Transforming: reconfiguring assets after entry

Transforming decides whether a route a firm has committed to becomes a lasting position. A foothold matters only if the firm can rearrange its assets after entry, and each route makes that harder in its own way. Internal ventures need support from the parent and a grace period before they are judged. Acquired units need to be integrated, and they send learning back to the parent. Alliances need constant coordination and governance that keeps shared knowledge useful (Cohen and Levinthal, 1990; Lane and Lubatkin, 1998).

Whatever the route, how much a firm actually learns is limited by its *absorptive capacity*: its ability, which depends on what it already knows, to turn new investment and incoming knowledge into usable capability (Cohen and Levinthal, 1990). A firm gets value from a partner or a target only as far as its related knowledge allows.

Transformation in the model is a continuing process. After entry, capability levels keep changing together, the parent cushions young ventures, alliance and acquisition spillovers move capability in both directions, and units that cannot reach scale are removed. This design gives two expectations. First, because the burden of transforming differs across routes, the route a firm took should leave a lasting mark on its survival. Second, because transformation builds advantage through feedback, the spread in outcomes should be path-dependent, produced as the system runs rather than inherited from where firms started.

2.4 An evolutionary reading: routines, selection, and institutions

Read in the vocabulary of evolutionary economics, the model is a population of routine-following firms under market selection (Nelson and Winter, 1982). The decision rules of Section 3 are routines

in Nelson and Winter's sense: every incumbent applies the same fixed procedure, an aspiration-triggered search followed by a noisy choice over a shared valuation, and firms differ in the state the procedure reads, not in the procedure itself. Dawid and Harting (2012) discuss how firm behavior of this kind is captured in agent-based models of industry evolution. Selection works through two channels. Demand is allocated with increasing returns that concentrate share on the firms with the strongest position (Arthur, 1989; Metcalfe, 1994), and units that fall below a traction floor exit; together these produce the shakeout and the concentration of the industry life cycle (Klepper, 1996, 1997; Klepper and Simons, 2000). Competition and collaboration are two faces of the sourcing decision: an acquisition removes a target that a rival could have bought, while an alliance creates a joint unit and leaves its partner in the market, as in simulation studies of innovation networks (Gilbert et al., 2001). Models of R&D collaboration add a result relevant here: stable collaboration structures need not be efficient (Dawid and Hellmann, 2014). Institutions enter in two places. The contracting regime, summarized by how contractible the needed capability is, decides whether a partner accepts an alliance and how much owning a tacit capability is worth (Williamson, 1985), and in a robustness experiment a merger-review-type friction raises the price of acquisition as the dominant design consolidates. The design follows agent-based models of innovation and industry dynamics (Silverberg et al., 1988; Winter et al., 2003; Dawid, 2006; Dosi et al., 2010), including models of capability accumulation (Gräbner and Hornykewycz, 2022), and history-friendly models that read simulated industries against documented ones (Malerba et al., 1999; Fontana and Zirulia, 2015; Malerba et al., 2016). Closest to our setting, agent-based work shows that a change of technological regime raises productivity heterogeneity and market concentration (Dawid and Hepp, 2022); we ask how incumbents source capabilities across such a regime change.

2.5 Interactions among the stages, and what an agent-based model represents

Sensing, seizing, and transforming are distinct ideas, but disruption response runs through their constant interaction, and that interaction forms a loop. Sensing decides who searches, seizing decides which routes are attempted, and transforming decides which affiliates survive and how fast entry barriers rise. Those outcomes then change what every remaining incumbent can sense, and can afford to seize, the next year.

These interactions also cross firm boundaries, because all the firms share one competitive environment. An acquisition leaves fewer targets for everyone else, an alliance reshapes the web of ties, and an early mover raises the barrier later movers must clear. The value of a firm's response therefore depends not just on its own actions but on how those actions meet its rivals'.

Three features of the problem shape the choice of method. Route choice is repeated, so the pool of targets runs down and constrains later choices; firms differ in resources, so no representative firm stands for the population; and the transition has no resting point during the period that matters. An agent-based model represents all three directly. It lets many different, boundedly rational firms interact under feedback and path dependence, so that market-level patterns emerge from individual choices instead of being assumed (Bonabeau, 2002; Harrison et al., 2007; Rand and Rust, 2011; Wu et al., 2024; Robertson, 2019).

The expectations above can be stated as four directional propositions that the simulation is designed to test and could refute.

P1: The most common route depends on the choices already made. As rivals' acquisitions use up the pool of targets, the mix of routes firms take shifts through feasibility, not

through a change in preference.

P2: Early entry gives an advantage that is a time-in-market effect, not an artifact of stronger or more-pressured firms moving first, and that advantage erodes as the transition speeds up and makes an early lead obsolete.

P3: Continuous resources sort firms into routes, but the size of the resulting spread in outcomes is path-dependent, produced as the system runs rather than inherited from the firms' starting positions.

P4: The boundaries between the mode-dominant regions are set by a few basic features of route economics, with each mode's region governed mainly by its own cost together with how contractible the needed capability is.

The propositions map one-to-one onto the four subquestions. Each construct enters the model as a specific operation. Sensing is the attention gate and search pressure of Eqs. 10–11, which decide whether a firm evaluates entry in a given year. Seizing is the screening, valuation, feasibility, and choice rules of Eqs. 12–17 together with the execution rule of Eq. 18. Transforming is the learning rule of Eq. 20 and the support and patience rules of Section 3.5. Each proposition is a theoretical expectation; the experiment that examines it is a separate object, named in Table 1 and specified in Section 3.7. Section 3 specifies the model that tests them.

Table 1: From theory to model. Each dynamic capability in Section 2, and the cross-firm feedback that links them, becomes a mechanism in the model, generates one or more of the propositions, and is identified by a specific test in Section 4.

Capability	Theoretical claim	Model operationalization	Prop.	Identifying test
Sensing (attention)	Performance feedback drives search; inertia holds attention on the legacy core	Sigmoid attention gate on entry: rises with performance gap, share loss, frontier gap, disruption, prior entry; falls with customer dependence, age, low adaptability	P1	Feasibility vs. preference split (SQ1)
Seizing: route	Build, partner, and acquire trade off speed, control, and access; the fit is contingent	Softmax over one shared payoff kernel minus a route cost, with symmetric frictions	P1, P4	Depletion knockout (K1); cost and contractibility sweeps; phase map
Seizing: timing	Early commitment can win a first-mover advantage that decays as a lead obsolesces	Entry year set by the attention gate; transition pace swept across six levels	P2	K3 randomized timing; K5 randomized entry order
Transforming	Absorptive capacity and route-specific integration burdens gate a durable position	Post-entry capability co-evolution, parent support, two-way spillovers, sub-scale exit floor	P3	K4 homogenize resources; survival analysis
Cross-firm feedback	One firm's move changes the next firm's options; increasing returns raise barriers	Depletion of the target pool by acquisition; Arthur sharpening; endogenous entry barrier	P1, P4	Depletion (K1); sharpening-off; decoupled-HHI check

3 The model: elements, rules, and what they stand for

The model is a computational laboratory for one specific problem. A population of legacy incumbents, the established firms that hold the declining business, must move from that shrinking legacy market into a linked, growing emerging market while the technology frontier keeps moving. Each year, every incumbent decides whether to build, partner, acquire, or wait. We do not fit the model to data. Instead, we build it so that which mode comes out ahead is something the simulated population produces.

Three design principles run through every part of the model, and we state them once here because they justify most of the choices that follow. First, *headline patterns must be emergent*: anything we later report as a finding, such as concentration, the shakeout, or the mode-dominant regions, has to be produced by a rule we can point to and switch off, not set by a parameter. Second, *no route may be privileged*: the three entry routes are scored by one shared payoff formula and every route is held to the same outcome thresholds, so that any "one mode wins" result comes from the payoffs rather than from how outcomes are counted; where the routes do differ, in costs, starting positions, transfers, dividends, and patience windows, the difference is stated with its value. Third, *every candidate mechanism must be switchable*, which is what makes the knockout tests of Section 4 possible. Where a design choice below follows from one of these principles, we say which.

The section is organized in the order a reader needs to follow one year of the model. Section 3.1 states what the model contains and the six-step annual schedule, with one process figure. Section 3.2 specifies the agents, their state, and how the population is initialized, and closes with an inventory of every element and a glossary. Section 3.3 specifies the two markets: how demand moves between them, how the new regime locks in, and how demand is allocated among firms. Section 3.4 specifies the incumbent's decision, from attention to execution, and the affiliate each route creates. Section 3.5 specifies what happens after entry: operations, learning, transfers, and exit. Section 3.6 defines every outcome measure and maps each phenomenon the paper studies to the rules that generate it and the test that identifies them. Section 3.7 sets out the experiments. Appendix A restates the model in the ODD+D protocol format (Grimm et al., 2010; Müller et al., 2013; Railsback and Grimm, 2019), Appendix B lists every governed parameter with its provenance and range, and Appendix C pairs each design choice with its evidence; the main text states the model's core mechanisms and procedures, and the appendices add provenance and supporting detail. Stating every rule and coefficient in the body follows calls for complete, reproducible documentation of agent-based models (Dawid et al., 2019).

Not every rule carries the same weight, and the inventory in Section 3.2 marks which is which. Six mechanisms are load-bearing, in the sense that a headline result depends on them and a knockout or a sweep tests them: the competitive-substitution rule that moves demand between the two markets (Eq. 3); the sharpened allocation of demand with its exit floor (Eq. 8); the attention gate (Eq. 10); the symmetric valuation and costing of the four routes (Eqs. 14–17); the depletion of the target pool by acquisition together with bilateral acceptance of alliances (Eqs. 13 and 18); and learning gated by absorptive capacity (Eq. 20). The supporting rules for margins, financing, capacity adjustment, and inter-firm transfers (Eq. 19) are not tested by knockouts, but they are not idle: they carry each firm's cash and capacity from one year to the next and therefore decide whether a young affiliate outlasts its gestation window. Their coefficients are held at baseline values throughout, with one exception: the scale-economy exponent in Eq. 19 is among the parameters the fifty-percent perturbation of Section 4.8 varies. Baseline coefficient values are given in the text next to each equation; Appendix B records the source and governed range of every governed one

and marks those varied in a reported analysis. Throughout, $\sigma(x) = 1/(1 + e^{-x})$ is the logistic function, and $\text{clip}(x)$ without stated bounds truncates x to $[0, 1]$; other bounds are written out.

3.1 Model overview and annual schedule

The model contains two markets and three kinds of actor. The *legacy market* is the declining business; the *emerging market* is the one the incumbents must enter. *Legacy incumbents* are the decision-makers. *Emerging independents* are the smaller firms that arise in the new market on their own; each is at once a rival, a possible acquisition target, and a possible alliance partner. An *affiliate* is the unit an incumbent creates when it enters: an internal build, an acquired unit, or an alliance unit held at an equity stake. The interaction that ties the population together runs through three channels. Incumbents draw on one finite pool of independents, and an acquisition removes its target from that pool for everyone who decides later. Demand in the emerging market is allocated with increasing returns, so the leaders pull away as the new technical standard settles. And the entry barrier rises with that consolidation, so an early affiliate raises the hurdle later movers face. One firm's move therefore changes the menu and the hurdle the next firm reads, which a one-shot choice model does not represent.

Time in the model moves in one-year steps. We use a yearly step because the decisions we study, entering, acquiring, and allying, run on roughly yearly planning cycles; a finer step would add detail without adding any new decisions. Main-experiment runs last sixty years (fifty in the sensitivity and phase-map runs of Experiment 3; Table 8 gives the horizons of the robustness designs), long enough for the emerging market to grow, peak, and shake out, so the long-run consequences of an early commitment show up within a single run. Within the one product category, demand is split between a legacy market that declines and an emerging market that grows as adoption speeds up and then settles. The firms are diverse: each has its own attributes and follows the same rules from a different starting point.

Each year executes the same six steps in a fixed order, and the order itself is a design choice. Figure 1 draws the schedule.

(1) *Market update.* The model advances the two markets: it draws the innovation shock, reads the frontiers and concentration off the population, updates the dominant-design strength, the emerging share, the entry barrier, the life-cycle stage, and the disruption flag (Section 3.3). We call the resulting bundle of conditions the regime.

(2) *Incumbent decisions.* Legacy incumbents, in random order, sense, screen, and seize (Section 3.4). A firm that acts creates an affiliate; its target is consumed or its partner engaged.

(3) *De novo entry.* New emerging independents arrive (Eq. 9).

(4) *Demand allocation.* Demand is handed out in the legacy market and then in the emerging one, in proportion to each active firm's competitive position, sharpened in the emerging market (Eq. 8).

(5) *Operations, learning, and transfers.* Each active firm and affiliate in turn settles its profit and finances, adjusts its capacity, and then updates its capability stocks with that year's research spending; once every firm has done so, capability and cash move between affiliates and parents (Section 3.5).

(6) *Aging and selection.* Ages and gestation timers advance, and firms and affiliates that cannot sustain themselves exit (Section 3.5).

Two features of this schedule need explaining. We run the market update before any firm decides, so that every firm reads the same regime within a year; and we randomize the order in which incumbents decide, so that no firm gets a built-in advantage from moving first. The second feature matters because entry timing is one of our subquestions (SQ2): any timing advantage we later find has to come from the firms' behavior, not from the order we happened to use.

Two kinds of process run through the schedule, and Figure 1 colors them apart. Sensing, seizing, and transforming (steps 2 and 5) are the dynamic capabilities the firm carries out, matching Section 2; screening is the front end of seizing, the step in which the firm picks out the specific partner or target it might commit to, kept separate so that candidate-finding and route choice can be tested on their own. The market update, de novo entry, demand allocation, and selection (steps 1, 3, 4, and 6) are processes of the market, not of any firm; selection is the survival filter of evolutionary economics and organizational ecology that decides which committed routes last (Jovanovic, 1982; Hannan and Freeman, 1984). The figure's boxes follow execution order; the grouping into capabilities and market processes is shown by color.

The schedule matters because it does not stop at the moment of choice. It carries each route forward into capability growth, financing pressure, barrier-building, and exit, so the long-run consequences of an early commitment appear within the run. All the randomness comes from one seeded generator, and every output records its seed, parameter values, and code version, so each number we report can be reproduced.

3.2 Agents, state, and initialization

The actors are those introduced in Section 3.1; this subsection specifies what each carries and how the population is initialized. For an alliance unit, the equity stake is the fraction of the joint unit the incumbent owns.

Every firm and affiliate carries four capability stocks, knowledge K_i, technology T_i, adaptability A_i, and legitimacy L_i, each on a zero-to-one scale, together with cash F_i, capacity C_i, captured demand q_i, market share s_i, and age a_i. Each also keeps the short histories the rules read (a three-year profit history for the aspiration gap and a four-year share history for the share-decline signal, both in Eq. 10) and, for affiliates, the gestation timers that shield a young unit (a build ramp, an integration window, or a coordination-drag window, plus a newness window; Section 3.5). We use these four stocks because they are the smallest set the theory of Section 2 needs: knowledge and technology capture what a firm can do, adaptability captures how readily it can change, and legitimacy captures whether the market accepts it. Table 3 lists the symbols and their domains.

How these stocks are set at the start decides whether the model can later separate sorting by resources from inequality that competition produces. Resources are drawn continuously within three overlapping seed size classes, and decision and payoff rules use the realized resources rather than the class labels; a model that assigned firms to fixed buckets and read those labels back out of the results would have assumed the spread in outcomes instead of finding it. We draw each firm from a single underlying maturity factor $Z_i \sim \mathcal{N}(0,1)$ whose effects have evidence-based signs, so that attributes vary continuously and are correlated, with overlapping distributions across firms. For attribute α with signed loading ℓ_α, the standardized draw and the realized attribute are

$$X_{\alpha i} = \delta\,\ell_\alpha\,\zeta_i + \rho\,\ell_\alpha\,Z_i + \sqrt{1-\rho^2\ell_\alpha^2}\;\epsilon_{\alpha i}, \qquad \alpha_i = \mathrm{clip}_\alpha(\mu_\alpha + \nu\,\sigma_\alpha\,X_{\alpha i}), \tag{1}$$

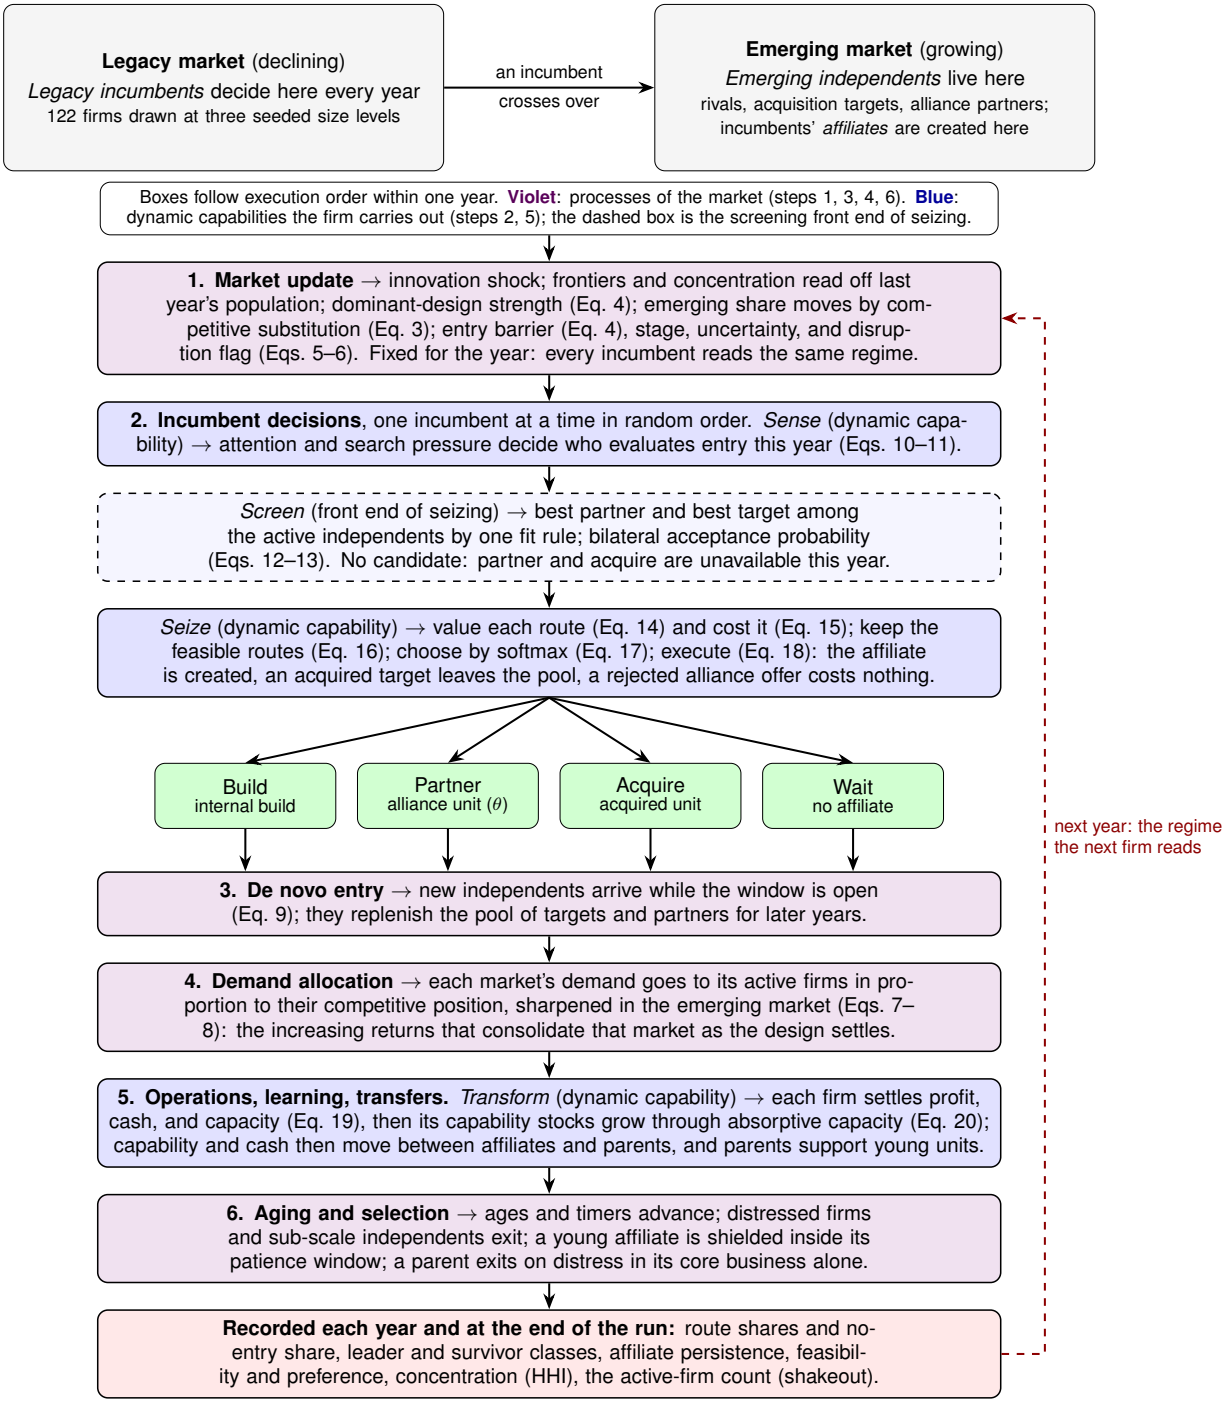


Figure 1: The model as one annual schedule, read top to bottom in execution order. Violet boxes are market processes (steps 1, 3, 4, and 6); blue boxes are the dynamic capabilities the firm carries out (steps 2 and 5), with the dashed box the screening front end of seizing. Equation numbers refer to the rules in Sections 3.3 to 3.5. The dashed red arrow is the cross-firm feedback: an acquisition removes a target from the pool, affiliates raise the entry barrier, and demand allocated with increasing returns consolidates the market, so the conditions the next firm reads are those the population has just produced.

where $\epsilon_{\alpha i}$ is independent standard noise, $\rho = 0.6$ sets the strength of the correlation, $\nu = 1.0$ scales the spread, and $\zeta_i \in \{-1, 0, +1\}$ with separation $\delta = 0.8$ places each incumbent's mean at a small, medium, or large seed level. The truncation clip_α is specific to each attribute: knowledge, technology, adaptability, and legitimacy are truncated to $[0, 1]$; age has a floor of one year and is rounded to whole years; cash has a floor of 1 for incumbents and 0.5 for emerging independents; and capacity has the floor $C_{\min} = 0.12$ and no upper bound. Emerging independents are drawn from the same equation with their own marginals and $\zeta_i = 0$. Because the loadings are scaled by ρ and the noise by its complement, within a size class every attribute keeps unit variance whatever ρ is, and any two attributes correlate at $\rho^2 \ell_\alpha \ell_\beta$. Setting $\rho = 0$ would remove the within-class correlation (and with $\delta = 0$ make the draws fully independent); the reported runs keep $\rho = 0.6$, and the K4 knockout collapses every firm, entrants included, to the attribute means. Table 2 gives the means, standard deviations, and loadings. Knowledge, technology, cash, capacity, and legitimacy load positively on maturity, in line with the view that competitive position is a stock built over time rather than bought at once (Penrose, 1959; Dierickx and Cool, 1989); adaptability loads negatively, in line with structural inertia (Hannan and Freeman, 1984). After setup, each incumbent is given a descriptive resource rank, the sum of the standardized values of $\log(1 + F)$, K, T, L, and $\log(1 + C)$, cut into terciles, that we use only for reporting and as a control in the survival analysis; no decision or payoff rule ever reads it. This lets us separate sorting from starting endowment later.

Table 2: Initialization marginals by class: mean and SD for legacy incumbents and emerging independents, with the signed latent-factor loading of Eq. 1. The loading column gives the legacy magnitudes; emerging independents share the signs with smaller magnitudes (age +0.50, F +0.50, K +0.40, T +0.30, A −0.30, L +0.40, C +0.50). Incumbents start older, larger, and more legitimate, whereas independents start smaller and cash-thinner but younger, more adaptable, and ahead on the emerging technology frontier. Realized values are truncated by the attribute-specific bounds stated with Eq. 1.

	Legacy incumbent		Emerging independent		
Attribute	Mean	SD	Mean	SD	Loading
Age	14.0	7.0	4.0	2.0	+0.85
Financial resources F	18.0	11.0	7.0	3.0	+0.60
Knowledge K	0.52	0.15	0.60	0.14	+0.55
Technology T	0.45	0.12	0.62	0.12	+0.35
Adaptability A	0.52	0.16	0.72	0.14	−0.55
Legitimacy L	0.55	0.18	0.30	0.10	+0.65
Capacity C	1.30	0.80	0.55	0.25	+0.60

The two classes of firm therefore enter at different scales, and a firm's scale is what later shapes its options. A legacy incumbent starts large and cash-rich, with more than twice the production capacity and the cash of a typical emerging independent, though it is older, less adaptable, and behind on the emerging frontier. An emerging independent starts small and cash-thin, but younger, more adaptable, and ahead on the new technology. Within each class, the draws are continuous and overlapping, not sorted into tiers.

We seed the incumbents across the three size levels in a right-skewed mix, with more small firms than large: fourteen large, thirty-six medium, and seventy-two small. These counts, and the separation between the levels, are illustrative normalization choices, and Section 4.8 shows that the results

do not depend on them. The size-class label is never an input to any decision or payoff rule; we use it only to set this mix and, after the run, as a control in the survival analysis. On any single attribute the three classes overlap heavily: on most of them about one small firm in six sits above the median large firm. On the combined resource index used in the sorting analysis they separate more sharply. The classes are distinguishable but not destiny: the rules always read the continuous resource, never the discrete class. The emerging market starts with six to ten independents drawn from the emerging-independent columns of Table 2 at the middle level $\zeta = 0$, and later entrants are drawn from the same distribution.

The three seeded size levels are therefore a way of setting where each incumbent's draw is centered; the reporting terciles are a different object, computed afterwards from the realized draws and used only to describe results. Size enters the decision through the firm's cash, not through any size-specific rule, along two channels that Eqs. 16 and 17 make precise. First, a route is feasible only if the firm can afford it, and because the price of an acquisition rises with the size of the target, a small, cash-thin incumbent cannot acquire a target many times its size: that route is closed to it, while a large incumbent can take it. Second, the choice rule weighs a route's cost against the firm's own cash, so the same price deters a small firm more than a large one. These two channels are the mechanism behind the resource sorting in Section 4.4, where richer firms enter more often by every route, most of all by building, while leaner firms more often stay out.

Table 3 collects every element of the model in one inventory, grouped into arenas, actors, firm state, market state, rules in execution order, and the levers the experiments vary. Each element is paired with the real-world object or phenomenon it stands for and with the equation that defines it, and the load-bearing rules are marked. Table 4 is a plain-language glossary of the special terms, for readers who meet them for the first time.

Table 3: Model elements and what they stand for. Every entity, state variable, rule, and lever in the model, the formal object that represents it, the real-world referent or phenomenon it stands for, and where it is defined. Rules marked † are the load-bearing mechanisms named in the text; unmarked rules are supporting rules (entry, operations and finance, selection, scoring), one of which also carries a knockout (the share floor in selection).

Element	Formal object	What it stands for	Defined in
Arenas and time			
Legacy market L	Demand D_t^L; active-firm set	The declining business the incumbents hold	Eq. 2
Emerging market E	Demand D_t^E; active-firm set	The new business the incumbents must enter	Eq. 2
Year t; horizon T	$t = 0, \ldots, 60$	Annual planning cycles over one full industry life cycle	Sec. 3.1
Actors			
Legacy incumbent i	Firm with state $(a_i, F_i, K_i, T_i, A_i, L_i, C_i, q_i, s_i)$	The established firm that must choose build, partner, acquire, or wait	Sec. 3.2
Emerging independent c	Same state vector, drawn smaller, younger, and ahead on T	The de novo entrant that is at once a rival, an acquisition target, and an alliance partner	Sec. 3.2; Eq. 9

Table 3 (continued)

Element	Formal object	What it stands for	Defined in
Affiliate	Internal build, acquired unit, or alliance unit at equity stake θ	The unit through which an incumbent operates in the new market	Eq. 18
Firm state			
Knowledge, technology, adaptability, legitimacy	$K_i, T_i, A_i, L_i \in [0,1]$	What the firm can do (K,T), how readily it changes (A), whether the market accepts it (L)	Sec. 3.2
Quality (derived)	$Q_i \in [0,1]$	The weighted blend of the four stocks that customers, partners, and investors respond to	Eq. 7
Financial resources	$F_i \in \mathbb{R}$	Cash and financing capacity; the firm's slack and what it can afford	Eq. 19
Capacity	$C_i \in \mathbb{R}_+$	Scale of operations	Eq. 19
Captured demand; share	q_i; $s_i \in [0,1]$	The position the firm has won in its market	Eq. 8
Age	a_i	Inertia; older firms attend less to the new market	Eq. 10
Gestation timers	Ramp, integration, coordination drag, newness window	The period in which a young affiliate is still ramping and is shielded from exit	Sec. 3.5
Market state (the regime every firm reads)			
Emerging share	$m_t \in [0,1]$	How far the transition has gone	Eq. 3
Technology frontiers	$T_t^L, T_t^E = \max_i T_i$ over active firms	The best technology each regime has reached; produced by the population, not imposed	Eq. 3
Fitness gap	FitGap_t (signed)	The relative merit of the new regime over the old	Eq. 3
Dominant-design strength	$S_t \in [0,1]$	How settled the new technical standard is	Eq. 4
Entry barrier	$B_t \in [0,1]$	The hurdle a newcomer must clear	Eq. 4
Concentration	HHI_t^E	How consolidated the emerging market is	Eq. 4
Stage and uncertainty	Life-cycle stage; u_t	How unresolved the emerging market still is; the source of option value	Eq. 5
Disruption flag	Binary	The point at which the regime has tipped	Eq. 6
Innovation shock	ε_t	Unanticipated jumps in the new technology	Eq. 3
Rules, in the order the loop runs them			
Market update†	Eqs. 2–6	Competitive substitution between regimes and lock-in of the new one	Sec. 3.3
Sense†	Attention att_i; search pressure P_i	Who searches: aspiration-triggered search held back by inertia	Eqs. 10–11
Screen†	Fit; bilateral acceptance	Finding a partner or target, and the counterparty's consent	Eqs. 12–13
Seize†	Route values V_r, costs, feasibility, softmax	The build, partner, acquire, or wait decision	Eqs. 14–17

Table 3 (continued)

Element	Formal object	What it stands for	Defined in
Affiliate creation†	Initial state by route; depletion of the target pool	What each route delivers on the day of entry	Eq. 18
De novo entry	Poisson arrivals	The window of opportunity for new firms; the supply of targets and partners	Eq. 9
Demand allocation†	Share $\propto \text{pos}_i^{\beta}$ (emerging; $\beta = 1$ in legacy)	Increasing returns: the leader pulls away as the design settles	Eq. 8
Operations and finance	Profit, cash, capacity updates	Bookkeeping that carries a position forward	Eq. 19
Transform†	Learning gated by absorptive capacity	Post-entry capability growth and its route-specific frictions	Eq. 20
Select	Exit rules by entity and route	Market selection; the liability of newness	Sec. 3.5
Outcome classification	Leader, survivor, failed, no entry	The end-of-run scorecard, applied identically across routes	Sec. 3.6
Levers varied in the experiments			
Pace	Substitution rate a	How fast the new regime can displace the old	Eq. 3
Target liquidity	Entry rate ρ_{entry}	How quickly new targets and partners appear	Eq. 9
Contractibility	$\chi \in [0, 1]$	How far the capability can be traded at arm's length rather than owned	Eqs. 13, 14
Route costs	c_b, c_p; price multiple μ and friction ξ	The own cost of building, partnering, and acquiring	Eq. 15
Partner reliability	Acceptance center θ_{acc}	How readily a partner agrees	Eq. 13
Slack	Cash normalizer κ_F	How much financial headroom the firm feels it has	Eq. 11
Integration friction	Base integration loss ℓ_0	What an acquisition loses in integration	Eq. 14
Uncertainty	Stage values u^{stage}	How unresolved the market is at each stage	Eq. 5
Knockout switches	K1, K3, K4, K5 (a barrier-freezing K2 is implemented but not used); sharpening, exit floor, contractibility gains	Remove one mechanism at a time to test what it generates	Sec. 3.7

3.3 Market dynamics and demand allocation

The market update builds the environment the incumbents read each year: a disciplined set of opportunities, not an estimated demand system. Total category demand grows along a simple curve, and the legacy market's slice of it is the complement of the emerging share, with a small exogenous decline on top:

$$D_t = D_0(1 + g_c)^t, \qquad D_t^L = \text{clip}\big((1 - m_t)(1 - \delta_L)^t,\, 0.05,\, 0.98\big)\, D_t, \qquad D_t^E = D_t - D_t^L, \tag{2}$$

Table 4: Glossary of terms. Each term is given a plain-language meaning and the model quantity that represents it, with a pointer to where it is formalized.

Term	What it means	Model representation
Regime	The bundle of market conditions every firm reads in a given year	Frontiers, demand split, S_t, B_t from the market update (Sec. 3.3)
Technology frontier	The best technology so far achieved in a market	T_t^L, T_t^E: maximum technology level over active firms (Eq. 3)
Route (mode)	One of the four responses: build, partner, acquire, or wait	Choice set of the seize step (Eq. 17)
Affiliate	The entity an incumbent creates to enter	Internal build, acquired unit, or alliance unit at equity stake θ (Eq. 18)
Dominant-design strength	How settled the new technical standard is	Bounded index $S_t \in [0, 1]$ (Eq. 4)
Entry barrier	The obstacle a new entrant must clear	Bounded index $B_t \in [0, 1]$ (Eq. 4)
Disruption flag	A marker that the regime has tipped	Binary condition on share, frontier, and design (Eq. 6)
Contractibility	How far the needed capability is codified and tradable at arm's length, rather than tacit	Parameter $\chi \in [0, 1]$ acting on alliance acceptance and on the premium for owning (Eqs. 13, 14)
Target liquidity	How rapidly new acquisition targets and partners appear	Entry rate of emerging independents (Eq. 9)
Slack	A firm's financial headroom	Cash F_i relative to the scale κ_F in search pressure and to route costs in the choice rule (Eqs. 11, 17)
Attention	How strongly a firm's search is directed at the emerging market	Sigmoid index gating entry evaluation (Eq. 10)
Feasibility vs. preference	Whether a route is reachable for a firm vs. whether the firm wants it	Affordability-and-readiness rate vs. choice probability among feasible routes (Eq. 16; Sec. 3.6)
Absorptive capacity	The ability to turn investment and inbound spillovers into usable capability	Weighted function of K, A, L gating learning (Eq. 20)
Mechanism knockout	An experiment that disables one mechanism to test whether a pattern is generated by it	Switch parameters K1, K3, K4, K5 and ablations (Secs. 3.7, 4.8)
Phase diagram (map)	A chart of which mode wins over a plane of two conditions, where a mode wins a cell if it is the first route of the most incumbents that end as leaders (Sec. 3.6)	Area shares of the winning mode (Sec. 4.5)

with $D_0 = 100$, $g_c = 0.002$, and $\delta_L = 0.004$. We keep the total deliberately simple, because the strategic question is how demand moves *between* the two markets, and a simple total keeps that question separate from unrelated swings in overall demand.

Competitive substitution. The central design question is how the emerging share m_t grows, and our answer comes from a failure of the obvious alternative. An automatic diffusion curve would spread the new regime no matter what firms did, which guarantees the very disruption the model exists to study. We therefore use a competitive-substitution rule, a replicator law that extends Fisher–Pry logistic substitution: the emerging regime gains share at the legacy regime's expense in proportion to how fit it is by comparison (Fisher and Pry, 1971; Metcalfe, 1994; Strohmaier and Rainer, 2016):

$$\begin{aligned} m_{t+1} &= \mathrm{clip}_{[0,\,m_{\max}]}\big[\, m_t + a\, m_t(1-m_t)\,\mathrm{FitGap}_t \,\big], \\ \mathrm{FitGap}_t &= b\,(T_t^E - T_t^L) + c\,(S_t - \bar{S})_+ + \varepsilon_t, \qquad T_t^X = \max_{i \text{ active in } X} T_i. \end{aligned} \tag{3}$$

The substitution rate a (baseline 3.4) is the model's *pace* lever: it sets how fast a fitter regime can displace the incumbent one, and the experiments sweep it from a slow to a fast transition. The gap weights are $b = 1.0$ and $c = 0.3$, the consolidation floor is $\bar{S} = 0.40$, and the innovation shock ε_t fires with probability 0.05 at a magnitude of 0.06 scaled by a uniform draw on $[0.7, 1.3]$. The share starts at $m_0 = 0.06$ and is capped at $m_{\max} = 0.92$, so a residual legacy niche always remains. The implementation also allows an intrinsic-advantage term in the fitness gap, set to zero at baseline and in every experiment, so the emerging market must gain share through the frontier or the dominant design rather than through an autonomous push. The logistic factor $m_t(1 - m_t)$ makes the annual movement largest in mid-transition.

Two consequences of this rule matter for everything that follows. First, the frontiers are not imposed: each is the highest technology level reached by any active firm in its market, so the gap is produced by the population itself. Second, because the driver is a *signed* gap, a legacy frontier that is still ahead makes fitness negative, and the emerging share shrinks. The transition can succeed, stall, or fail depending on the starting frontiers and the pace, so disruption is a condition the experiments can turn up or down.

Lock-in. Once share is driven by fitness, we track how locked-in the regime is with three bounded indices. Concentration is the Herfindahl index HHI_t^E, the sum of squared demand shares among active emerging firms, which approaches one as a single firm comes to dominate and falls toward zero when many firms share the market. Dominant-design strength S_t, an index of how settled the new technical standard is, and the entry barrier B_t, the height of the obstacle a new entrant must clear, are bounded weighted sums of population state:

$$S_t = \mathrm{clip}\big(\omega_1 m_t + \omega_2 T_t^E + \omega_3\,\mathrm{HHI}_t^E\big), \qquad B_t = \mathrm{clip}\big(\eta_1 m_t + \eta_2 \bar{Q}_t^{\mathrm{ib}} + \eta_3\,\mathrm{HHI}_t^E\big), \tag{4}$$

with both weight vectors equal to $(0.35, 0.35, 0.30)$ and $\bar{Q}_t^{\mathrm{ib}}$ the mean of $0.40\,T_i + 0.35\,K_i + 0.25\,L_i$ over active emerging firms, an installed-base quality proxy that leaves capacity out. One timing detail matters for reproduction: S_t is computed at the start of the market update, before the share moves, so it reads the previous year's share and concentration, whereas the barrier, the stage, and the disruption flag read the updated share. The design-strength index captures the idea that a regime settles once it has scale, technical coherence, and consolidation (Anderson and Tushman, 1990; Suarez and Utterback, 1995); the barrier captures the way increasing returns and lock-in raise the hurdle later entrants must clear (Arthur, 1989).

The life-cycle stage is read from the emerging share (emergence below 0.15, growth below 0.45, maturity below 0.75, late thereafter), and the stage sets how unresolved the market is:

$$u_t = \text{clip}\big(u^{\text{stage}} + 0.3\,\varepsilon_t,\ 0.05,\ 0.55\big), \qquad u^{\text{stage}} = (0.30,\ 0.22,\ 0.14,\ 0.12) \text{ across the four stages,} \tag{5}$$

so that uncertainty is high while the category is forming and falls as it matures, with a shock adding to it. This u_t is what gives the flexible routes their option value in Eq. 14. Finally, a disruption flag, a yes-or-no marker that the regime has tipped, turns on when

$$m_t \geq m_D \ \wedge\ \Big[\, T_t^E \geq \max(\theta_F,\ T_t^L + \Delta_F) \ \vee\ S_t \geq \theta_S \,\Big], \tag{6}$$

with adoption gate $m_D = 0.30$, frontier gate $\theta_F = 0.65$, margin $\Delta_F = 0.08$, and design gate $\theta_S = 0.45$. The flag therefore fires only when the emerging share is large enough and either the emerging frontier breaks past the legacy frontier or a dominant design consolidates, never merely because a new technology exists.

Increasing returns and the shakeout. Concentration and the post-peak shakeout do not need separate machinery. We produce both from a single increasing-returns mechanism: the self-reinforcing advantage by which a firm that gets ahead gets further ahead (Arthur, 1989). Each active firm holds a competitive position, built from the same formula that later scores the entry routes:

$$\text{pos}_i \ =\ Q_i\,\sqrt{C_i}, \qquad Q_i \ =\ w_T\,T_i + w_K\,K_i + w_A\,A_i + w_L\,L_i. \tag{7}$$

Quality Q_i is a weighted score that blends the four capability stocks, with weights $w_T = 0.40$, $w_K = 0.30$, $w_A = 0.10$, and $w_L = 0.20$, and the square root of capacity builds in diminishing returns to scale. This form rests on two pieces of evidence. Treating demand capture as quality times scale follows the industry-evolution finding that survival needs both a competitive product and the capacity to serve demand (Klepper, 1996); the square-root term reflects the classic argument that the advantages of size grow less than proportionally (Penrose, 1959). Emerging demand is then handed out in proportion to a *sharpened* position, while legacy demand is handed out in proportion to $\text{pos}_i\,\text{fr}_i$, that is, with $\beta = 1$:

$$\text{share of } D_t^E \text{ to firm } i \ \propto\ (\text{pos}_i\,\text{fr}_i)^{\beta}, \qquad \beta = \text{clip}\big(1 + \gamma\,S_t,\ 1,\ \beta_{\max}\big), \qquad \gamma = 7.0,\ \beta_{\max} = 6.0, \tag{8}$$

distributed over six capacity-constrained rounds, so that no firm serves more than its capacity, with a final pass in which a quarter of any residual demand is served on backorder. The factor fr_i is the gestation friction: a young affiliate's position is discounted by 0.15 while its build ramp, integration window, or coordination-drag window runs and by a further 0.12 while its newness window runs, and $\text{fr}_i = 1$ otherwise. The penalty is the same for every route; only the window lengths are route-specific parameters, and at baseline they are equal. The exponent β rises with dominant-design strength. As the design settles, the leaders pull away, and firms that stay too small fall below the exit floor of Section 3.5. So the sharpening drives both rising concentration and the shakeout, the sharp post-peak drop in the number of active firms that the industry life cycle records (Klepper, 1996, 1997), and the exit floor is what turns falling share into exit. We cap the exponent so the market settles into an oligopoly, not a single winner. Section 4.8 tests both links: removing the sharpening removes concentration, weakening it shrinks the shakeout, and removing the floor removes the shakeout.

De novo entry. New independents arrive each year by a stage-modulated Poisson process:

$$k_t \sim \text{Poisson}(\lambda_t), \quad \lambda_t = \max\big(0,\ \rho_{\text{entry}}\,\mu_{\text{stage}}\,w_t\big), \quad w_t = \max\big(0.08,\ (1 - B_t)(1 - S_t)\big), \tag{9}$$

with base rate $\rho_{\text{entry}} = 3.0$, the model's *target liquidity* lever, a stage multiplier μ_{stage} of 1.10, 0.95, 0.55, and 0.18 from emergence to late, and an opportunity window w_t that closes as the barrier and the dominant design rise. Each drawn entrant c is admitted only if its quality clears half the barrier, $Q_c \geq 0.5\, B_t$. This is the window-of-opportunity pattern of industry evolution (Suarez et al., 2015; Klepper, 1996), and because every entrant is also a potential target or partner, the entry rate governs how fast the pool the incumbents draw on is replenished.

When each market quantity updates. Grouping the market rules here does not mean they execute at one point in the year, so we state the timing. In step 1 of the schedule the model draws the innovation shock, reads the frontiers and concentration off the population as it stood at the end of the previous year, computes the dominant-design strength S_t from the previous year's share and concentration, moves the emerging share by Eq. 3, and then computes the barrier, stage, uncertainty, and disruption flag from the updated share. These quantities are then fixed for the year, so every incumbent that decides in step 2 reads the same regime. De novo entry (Eq. 9) is step 3, after the incumbents have decided, so entrants replenish the pool for the following year. Demand allocation (Eqs. 7 and 8) is step 4 and uses each firm's capability stocks and capacity as they stand after that year's entries but before that year's learning. The frontiers move through the learning of step 5 and as firms enter and exit, not through a directly, which is why the pace lever a affects the frontier only indirectly (Section 3.5).

3.4 Incumbent decisions and affiliate creation

The behavioral core turns the dynamic capabilities of Section 2 into a linked sequence of steps. For each step we give the rule, its rationale, and, where there is an obvious alternative, the reason we did not use it.

Sense. A firm's attention to the emerging market is a sigmoid of an evidence-weighted sum of signals, five that raise it and three that lower it:

$$\begin{aligned} \text{att}_i &= \sigma\big(g_A(W_i - \theta_A)\big), \\ W_i &= \beta_1 g_i^{\text{perf}} + \beta_2 d_i^{\text{share}} + \beta_3 g_i^{\text{front}} + \beta_4\, \mathbf{1}[\text{disrupt}] + \beta_5\, \mathbf{1}[\text{prior}] \\ &\quad - \beta_6\, \text{trap}_i - \beta_7\, \text{age}_i - \beta_8(1 - A_i), \end{aligned} \tag{10}$$

with steepness $g_A = 2.2$, offset $\theta_A = 0.25$, and weights β_1 through β_8 of 1.10, 0.70, 0.95, 0.75, 0.20, 0.70, 0.18, and 0.55. The signals that raise attention are the aspiration-relative performance gap $g_i^{\text{perf}} = \text{clip}\big((\bar{\pi}_i - \pi_i)/(|\bar{\pi}_i| + 1)\big)$, where $\bar{\pi}_i$ is a three-year rolling mean of profit; the recent loss of legacy share d_i^{share}, the drop from the firm's highest share over the last four years to its current share; the distance to the emerging frontier $g_i^{\text{front}} = \text{clip}\big(\max(0, T_t^E - T_i)\big)$; the disruption flag; and a prior-exposure indicator $\mathbf{1}[\text{prior}]$, which switches on the first time the parent executes an entry and stays on thereafter, even if that affiliate later exits. The behavioral theory of the firm links all of these to stronger search (Cyert and March, 1963; Greve, 2003). The forces that lower attention are the customer-power trap $\text{trap}_i = \text{clip}(0.6\, s_i + 0.4\, \text{clip}(\pi_i/12))$, which captures how a firm that is doing well with its existing customers keeps looking at them (Christensen and Bower, 1996; Gilbert, 2005); age inertia $\text{age}_i = \text{clip}(a_i/35)$; and low adaptability (Tripsas and Gavetti, 2000). The sigmoid bounds attention and gives it the threshold-like character of aspiration-triggered problemistic search. Attention then combines with slack and adaptability into a search pressure,

$$P_i = \text{clip}\big(0.45\, \text{att}_i + 0.35\, \text{U}(\varsigma_i) + 0.20\, A_i\big), \qquad \varsigma_i = \frac{F_i}{F_i + \kappa_F}, \qquad \text{U}(x) = \text{clip}\Big(1 - \big(\tfrac{x - 0.35}{0.35}\big)^2\Big), \tag{11}$$

where ς_i is the firm's slack, its cash relative to a normalizer $\kappa_F = 30$ that the slack experiments vary, and U is a downward parabola peaking at intermediate slack. Search effort is highest at intermediate slack and falls off when the firm is either too cash-constrained to act or too comfortable to bother, the inverted-U relationship documented by Bourgeois (1981) and Nohria and Gulati (1996). A firm evaluates entry when $P_i \geq P^* = 0.52$ or, below the threshold, with a small preemption probability $\text{clip}(0.15\,\text{att}_i + 0.05\,\mathbf{1}[\text{disrupt}])$; a firm that acted within the last three years does not re-evaluate. Attention acts as a gate on whether a firm considers entry at all, not as a preference among routes. The reason is the one from Section 2.1: if attention carried a taste for one route, it would blur the line between who searches and which route a searcher takes, and SQ1 depends on keeping the two apart.

Screen. A firm that searches finds its best partner and its best acquisition target among the emerging independents, using the same fit measure for both. For a candidate c, capability distance is $\text{dist}(i,c) = \frac{1}{3}(|K_i - K_c| + |T_i - T_c| + |A_i - A_c|)$, similarity is $1 - \text{dist}$, and complementarity is $\text{clip}\big(0.5\max(0, T_c - T_i) + 0.5\max(0, K_c - K_i)\big)$. The fit is their mean together with the candidate's quality,

$$\text{fit}(i,c) = \text{clip}\Big(\tfrac{1}{3}\big(\text{similarity} + \text{complementarity} + Q_c\big)\Big), \tag{12}$$

so that fit rewards a similar capability profile, complementary strengths, and high quality in the candidate. Rewarding similarity and complementarity together follows relative absorptive capacity (Lane and Lubatkin, 1998): a firm learns most from a partner that overlaps enough to be understood and differs enough to be worth learning from. Empirically, cognitive proximity is the strongest driver of both the formation and the output of collaborations (Werker et al., 2019), which is why partner fit enters alliance acceptance below. For a partnership the firm also estimates the chance the candidate agrees to ally,

$$\text{acc}(i,c) = \sigma\big(2.5\,(r_{ic} - \theta_{\text{acc}})\big), \qquad r_{ic} = 0.5\,\text{fit}(i,c) + 0.5\,\text{clip}(Q_i) - \psi\,\frac{|C_i - C_c|}{C_i + C_c} + g_\chi^{\text{p}}(\chi - 0.5), \tag{13}$$

with acceptance center $\theta_{\text{acc}} = 0.60$, the model's *partner reliability* lever, leakage weight $\psi = 0.30$, and contractibility gain $g_\chi^{\text{p}} = 1.2$. Acceptance rises with the partner's attractiveness and with the firm's own quality, and falls when a large size gap signals a risk of the smaller partner's knowledge leaking to the larger one, the appropriation hazard of the relational view (Dyer and Singh, 1998). The firm selects the partner that maximizes fit $\cdot$ acc and the unowned target that maximizes fit. Candidates are the independents active that year. An independent that has already been acquired is owned and is no longer a target, whereas one that is already in an alliance remains a candidate partner for other incumbents. If no candidate exists, the partner and acquire routes are unavailable that year: their values are set to minus infinity, so the choice in Eq. 17 is made among the remaining routes. We use one screening rule for both routes, because separate rules could hand one route a hidden advantage at this stage.

Contractibility $\chi \in [0,1]$ enters here for the first time. Contractibility is how far the needed capability is codified and can be obtained at arm's length, rather than tacit and dependent on control (Williamson, 1985; Capron and Mitchell, 2012). A capability a partner can deliver under contract is one a partner will agree to trade, so high contractibility raises acceptance in Eq. 13; a tacit capability has to be owned, so low contractibility adds a premium to the committed routes in Eq. 14. We make the lever act on acceptance rather than on the partner route's value, for a concrete reason. A lever that raised the partner route's value would push firms toward partnerships regardless of whether a counterparty would agree, so firms would choose to partner and then fail to execute. Acting on acceptance keeps the choice tied to what can be executed and gives a

monotone partner-share gradient that flattens out when we switch the mechanism off. Recent evidence points the same way: firms whose strategy emphasizes value appropriation rather than value creation prefer acquisitions to alliances (Mallapragada et al., 2025); reading that contrast as one of contractibility is our interpretation. At baseline $\chi = 0.5$, where both terms vanish.

Seize: what each route is worth. Each route's value is the competitive position it would deliver, computed from the same formula $\text{pos}(\cdot)$ of Eq. 7 that drives actual demand, adjusted for three things: how soon the position materializes, how much of it an irreversible commitment forfeits under uncertainty, and how much owning is worth when the capability is tacit. Let $\varrho(g) = h/(h+g)$ discount a route whose position takes g years of gestation, with planning horizon $h = 3$; let $\mathcal{I} = \text{clip}(1 - \phi_{\text{irr}} u_t,\, 0.2,\, 1)$ with $\phi_{\text{irr}} = 0.6$ be the real-options discount on irreversible commitment; and let $\mathcal{M} = \text{clip}\big(1 + g^{\text{i}}_{\chi} \max(0, 0.5 - \chi),\, 1,\, 1.8\big)$ with $g^{\text{i}}_{\chi} = 0.8$ be the premium for internalizing a tacit capability. Then

$$
\begin{aligned}
V_{\text{build}} &= \text{pos}(K_i, T_i, A_i, L_0,\, \kappa_b C_i) \cdot \varrho(g_b) \cdot \mathcal{I} \cdot \mathcal{M}, && L_0 = 0.28,\ \kappa_b = 0.22,\ g_b = 3, \\
V_{\text{acquire}} &= \text{pos}(\text{target}) \cdot (1 - \ell) \cdot \mathcal{I} \cdot \mathcal{M}, && \ell = \text{clip}\big(\ell_0(1 + \text{dist}(i,c)),\, 0,\, 0.9\big),\ \ell_0 = 0.20, \\
V_{\text{partner}} &= \theta\, \text{pos}(\text{best of both}) + w_o\, u_t (1 - S_t), && \theta = 0.5,\ w_o = 0.45, \\
V_{\text{wait}} &= u_0 + 0.05\, u_t + 0.30\, \varsigma_i, && u_0 = 0.30,
\end{aligned}
\tag{14}
$$

where "best of both" takes $\eta \max(\text{own}, \text{partner})$ for each capability stock, and capacity is $\eta\,(\kappa_b C_i + C_c)$, as in Eq. 18, with coordination factor $\eta = 0.80$. What differs across the routes is therefore the starting position each one delivers, and that difference is the trade-off between building and acquiring itself. Build inherits the parent's capability but starts at a fraction κ_b of its scale, with low legitimacy L_0 and a slow ramp g_b, because an internal venture starts small. Acquire obtains a target's full running position at once, minus an integration loss ℓ that grows with the distance of the capability, because an acquisition delivers a ready position, at a price. Partner combines the best of both partners at the equity stake θ, trimmed by the coordination factor, and carries an option value that rises with uncertainty and falls as the design settles. Wait keeps the cash and the option to act later. Only the two committed routes carry the real-options discount $\mathcal{I}$ and the internalization premium $\mathcal{M}$; the flexible routes do not, because committing irreversibly gives up the value of waiting until uncertainty clears (Dixit and Pindyck, 1994; Folta, 1998; Kogut, 1991). Each starting position is set by a stated parameter, and the four that matter most for the results, the build scale κ_b, the integration loss ℓ_0, the equity stake θ, and the coordination factor η, are varied in the reported analyses.

Seize: what each route costs, and which are open. Route costs depend on the state of the market:

$$
\begin{aligned}
\text{cost}_{\text{build}} &= c_b(0.8 + 0.4 B_t), \qquad \text{cost}_{\text{partner}} = c_p(0.8 + 0.5 u_t), \\
\text{cost}_{\text{acquire}} &= \big(c_a + \mu\,(q_c + 0.5 C_c)\big)(1 + \xi\, S_t),
\end{aligned}
\tag{15}
$$

with $c_b = 9.5$, $c_p = 7.0$, $c_a = 5.0$, price multiple $\mu = 2.5$, and $\xi = 0$ at baseline. Building gets dearer as the barrier rises, partnering as uncertainty rises, and acquiring with the target's captured demand q_c and capacity, the practice by which acquisition prices scale with the target (Haspeslagh and Jemison, 1991); the friction ξ, which raises the price as the dominant design consolidates, is switched on only in the robustness runs that stand in for merger review. A route is open to a firm only if it is both affordable and ready:

$$
\text{feasible}_r \iff \text{cost}_r \le F_i \ \wedge\ \text{readiness}_r \ge \rho^*, \qquad \rho^* = 0.35, \tag{16}
$$

where readiness is $0.5K_i + 0.5A_i$ for build, the partner's fit for partner, and the target's fit for acquire; wait is always feasible. The threshold is the same for every route. This is the first of the two cash channels of Section 3.2: a cash-thin incumbent cannot buy a target many times its size.

Seize: the choice. Among the feasible routes the firm chooses by a softmax rule, a random-utility rule that turns scores into choice probabilities:

$$\Pr(r) \;=\; \frac{\exp\big((V_r - \lambda\, c_r)/\tau\big)}{\sum_{r' \text{ feasible}} \exp\big((V_{r'} - \lambda\, c_{r'})/\tau\big)}, \qquad c_r = \frac{\text{cost}_r}{F_i + \kappa}, \qquad \lambda = \kappa = 1,\ \tau = 0.33. \tag{17}$$

The cost ratio c_r is the second cash channel: the same price deters a small firm more than a large one. The temperature τ is what makes the firm boundedly rational: it does not optimize, but picks among routes with a probability that rises with their score (McFadden, 1974). As the temperature falls toward zero, the rule becomes deterministic and the firm always takes the single best-scoring route, a limit the robustness checks approach with $\tau = 0.05$. We choose this random-utility rule over the two obvious alternatives for reasons tied to our question. A best-response rule would assume the firm already knows the payoff of every route on a frontier that has not settled, which is the very uncertainty the model exists to study. A reinforcement-learning rule would add tuning parameters with no empirical target. Scoring routes by a grounded value and then choosing with some noise instead lets adoption waves emerge, without assuming firms have already solved the problem. The noise the rule allows also matches the behavioral premise that firms satisfice under uncertainty rather than optimize (Cyert and March, 1963; Nelson and Winter, 1982).

The affiliate a route creates. A chosen route instantiates an affiliate with an explicit initial state, and this is where the routes' expected starting positions become actual ones:

$$\begin{aligned}
\text{build:}\quad & (K, T, A) = (K_i, T_i, A_i),\ L = L_0,\ C = \max(C_{\min}, \kappa_b C_i);\\
\text{acquire:}\quad & \text{the target is converted in place, keeping its state, age, cash, and history,}\\
& \text{and leaves the pool;}\\
\text{partner:}\quad & (K, T, A, L) = \eta \max(\text{own}, \text{partner}),\ C = \max\left(C_{\min}, \eta\left(\kappa_b C_i + C_c\right)\right),\ \text{held at stake } \theta;
\end{aligned} \tag{18}$$

A build or alliance unit is seeded with cash equal to half its route cost; an acquired unit keeps the target's cash and receives the same injection on top. Each affiliate then starts the route's gestation timers, a build ramp, an integration window, or an alliance coordination-drag window (each three years at baseline), plus a two-year liability-of-newness window. Because an acquired unit keeps the target's age, its patience window in Section 3.5 runs from the target's founding, not from the acquisition. For the partner route, the acceptance probability of Eq. 13 is realized at execution: the alliance forms with probability acc, and a rejected offer costs nothing, creates no affiliate, does not start the three-year cooldown, and leaves the firm free to evaluate again the next year. Only acquisition depletes the pool: an acquired target is converted into the acquirer's affiliate and leaves the pool of independents, whereas an alliance partner stays available to other incumbents. This is the depletion channel that lets one firm's choice change the next firm's menu, and it is the mechanism the K1 knockout removes by copying the target instead of converting it. The copy starts as a new unit of age zero with a cash seed of half the route cost while the original independent stays active, so under K1 an acquired unit's patience window runs from the acquisition. An alliance unit's capacity draws on its partner's capacity while the partner stays active with its own, so an alliance adds capacity to the market instead of moving it.

What the routes share and where they differ. The routes share one valuation formula (Eq. 14), one set of outcome thresholds (Section 3.6), and one demand kernel: profit, cash, learning,

and demand all flow from Eq. 7 for every unit. They also share the size of the temporary gestation penalty, which applies during a build ramp, an acquisition's integration window, or an alliance's coordination-drag window. Alliance value is weighted by the explicit equity stake θ, so an alliance-heavy firm can still lead, and the coordination factor η trims the combined bundle an alliance delivers (Kale and Singh, 2007). The routes differ in their costs (Eq. 15), their starting positions (Eq. 18), and their post-entry rules, the transfer and dividend rates and patience windows of Section 3.5; those post-entry rates are held at their baseline values in every reported analysis.

A worked year. Table 5 traces one incumbent through step 2 of the schedule with every input shown. The values are illustrative, chosen to sit inside the baseline ranges, not logged from a run, and each row is computed from the equation it names. The table shows two things. Waiting is the most likely single draw for this firm in this year, which is common at baseline: entry is spread over several years of re-evaluation. And the calculation applies the same rules to every route except in what each route delivers (Eqs. 14 and 18), its readiness test (Eq. 16), and its cost (Eq. 15), which is the symmetry principle at work. The same firm with twice the cash would see its cost ratios halved, which favors the entry routes over waiting, but its search pressure would then fall to 0.49, below the evaluation threshold of 0.52, so in that year it would evaluate only through the small preemption probability of about 0.06.

Table 5: A worked year for one incumbent. Illustrative values inside the baseline ranges, not a logged run; each block is computed from the equation named, with the baseline coefficients. The realized decision is a single draw from the choice probabilities shown.

Step	Inputs and result
Inputs: the firm	$K = 0.55$, $T = 0.50$, $A = 0.50$, $L = 0.60$; cash $F = 20$; capacity $C = 1.4$; legacy share $s = 0.03$; profit $\pi = 1.0$ against a three-year mean of 1.6; highest share in the last four years 0.05; age 15; no prior entry
Inputs: the regime	emerging frontier $T^E = 0.70$; emerging share $m = 0.30$ (growth stage); uncertainty $u = 0.22$; design strength $S = 0.35$; barrier $B = 0.40$; disruption flag off; contractibility $\chi = 0.5$
Inputs: best independent	$K_c = 0.65$, $T_c = 0.72$, $A_c = 0.70$, $L_c = 0.35$; capacity $C_c = 0.6$; captured demand $q_c = 0.62$; the same firm is the best partner and the best target
Sense (Eqs. 10–11)	signals: performance gap 0.23, share decline 0.02, frontier gap 0.20, customer trap 0.05, age 0.43, inertia 0.50; $W = 0.07$; attention 0.40; slack 0.40; search pressure $0.62 \geq 0.52$, so the firm evaluates entry
Screen (Eqs. 12–13)	distance 0.17, similarity 0.83, complementarity 0.16, candidate quality 0.62; fit 0.54; acceptance 0.39
Value (Eq. 14)	build 0.11 (small unit, half realized within the horizon, discounted for irreversibility); partner 0.29; acquire 0.32 (integration loss 0.23); wait 0.43
Cost and feasibility (Eqs. 15–16)	in the order build, partner, acquire: costs 9.1, 6.4, 7.3, all below cash 20; readiness 0.53, 0.54, 0.54, all above 0.35; all four routes feasible; cost ratios 0.43, 0.30, 0.35, and 0 for wait
Choice (Eq. 17)	scores build -0.32, partner -0.01, acquire -0.03, wait 0.43; probabilities build 0.06, partner 0.16, acquire 0.15, wait 0.62
Execution (Eq. 18)	one route is drawn; a wait leaves the firm free to evaluate again next year; a partner draw forms the alliance with probability 0.39 and otherwise costs nothing; an acquire draw converts the target and removes it from the pool; a build draw creates a unit at 0.22 of the parent's capacity

3.5 Post-entry dynamics, operations, and exit

Operations and finance. Each active firm and affiliate first updates its accounts:

$$\begin{aligned}
\pi_i &= q_i \,\mathrm{mgn}_i - 0.18\sqrt{C_i} - 0.12\, q_i \,\mathrm{clip}\big((C_i/C^*)^{-e_s S_t},\, 0.4,\, 2.5\big),\\
\mathrm{mgn}_i &= \mathrm{clip}\big(0.28 + 0.16(Q_i - \bar{Q}_t^X),\, 0.08,\, 0.45\big),\\
F_i &\leftarrow F_i + \pi_i - \mathrm{rd}_i - \mathrm{capex}_i + \mathrm{ext}_i, \qquad \mathrm{rd}_i = 0.08\max(0, F_i), \quad \mathrm{capex}_i = 0.05\max(0, F_i),\\
C_i &\leftarrow \max\left(C_{\min},\, C_i + 0.28\left(\max(C_{\min}, q_i) - C_i\right)\right),
\end{aligned} \tag{19}$$

where the margin rises with the firm's quality relative to the mean quality $\bar{Q}_t^X$ of active firms in its own market, the fixed cost grows with the square root of capacity, and the unit cost carries a scale economy active only in the emerging market ($C^* = 1.0$, $e_s = 0.9$) that strengthens as the design settles. External finance ext_i arrives with probability $\mathrm{clip}\big(\sigma(1.8\,(Q_i - 0.5))\, n_i\big)$ and pays $4.5\,Q_i$ scaled by a uniform draw on $[0.3, 1]$, where Q_i is the firm's quality clipped to $[0, 1]$ and the need multiplier n_i is 1 when cash is thin ($F_i < 8$) and 0.35 otherwise, so investors back higher-quality firms more readily when internal cash runs low. Capacity adjusts partway toward demand each year and never falls below the floor $C_{\min} = 0.12$, which also bounds the capacity of a new build or alliance unit (Eq. 18).

Learning (transforming). After entry, each affiliate and its parent build up their capability stocks through investment, scaled by absorptive capacity $\mathrm{ac}_i = 0.45K_i + 0.35A_i + 0.20L_i$, the ability, which depends on what the firm already knows, to turn investment and incoming spillovers into usable capability (Cohen and Levinthal, 1990; Lane and Lubatkin, 1998; Zollo and Winter, 2002). With rd_i the research spending of Eq. 19 and a newness factor $\mathrm{nw} = 0.88$ applied while the newness timer runs, the stocks update, all clipped to $[0, 1]$, as

$$\begin{aligned}
K_i &\leftarrow K_i + e_K \,\mathrm{rd}_i \,\mathrm{ac}_i (1 - K_i)\,\mathrm{nw},\\
T_i &\leftarrow T_i + e_T \,\mathrm{rd}_i \,\mathrm{ac}_i (1 - T_i)\,\mathrm{nw} + \underbrace{\iota\, \max(0, T_t^X - T_i)\,\mathrm{clip}(1 - \delta_\iota g_t)}_{\text{imitation catch-up}} - \underbrace{\omega_0(1 + \omega_g g_t)\, T_i}_{\text{obsolescence}},\\
A_i &\leftarrow A_i + 0.05(1 - A_i)\,\mathrm{ac}_i, \qquad L_i \leftarrow L_i + 0.04(1 - L_i)\big(0.5 + 0.5\min(1, q_i)\big),
\end{aligned} \tag{20}$$

with $e_K = 0.11$, $e_T = 0.09$, base obsolescence $\omega_0 = 0.012$, and frontier progress $g_t = \max(0, T_t^E - T_{t-1}^E)$. We run investment through absorptive capacity instead of letting it raise capability directly, because the theory of Section 2.3 says a firm learns from a partner or a target only as far as its related knowledge allows; the gap-closing form $e \cdot \mathrm{rd} \cdot \mathrm{ac} \cdot (1 - X)$ encodes diminishing returns as a stock approaches its ceiling, the standard asset-accumulation specification (Dierickx and Cool, 1989). Legitimacy grows with demand served (q_i is captured demand), since a firm that customers buy from becomes accepted. Two first-mover channels sit inside the technology update: imitation catch-up (rate ι, damped by a moving frontier at $\delta_\iota = 12$), which lets followers close on the frontier when the pace is slow, and pace obsolescence (gain ω_g), which erodes a pioneer's early lead when the pace is fast (Suarez and Lanzolla, 2007; Lieberman and Montgomery, 1998). Both are switched off at baseline ($\iota = \omega_g = 0$) and turned on only for the timing experiment and the first-mover arm of the knockout panel. This limits the first-mover pace result to those runs.

Transfers and parent support. Inter-firm transfers then move capability and cash along route-specific channels, and every capability transfer only raises the lagging side. A build unit draws knowledge and technology toward its parent's levels at rate 0.18 and the parent learns back at 0.08; an acquired unit and its parent exchange knowledge in both directions at 0.10; an alliance unit passes knowledge and technology to both sponsors at 0.10. Acquired units remit half of any

positive profit to the parent and alliance units 0.30 to the incumbent sponsor and 0.20 to the partner, while build units remit nothing. A sponsor with cash above 8 supports a unit inside its patience window whose cash has fallen below 4: it tops the unit up to 4 plus 0.35 of the unit's loss, paying at most its own cash in excess of 8. The margin band, financing, and capacity-adjustment rules are accounting normalizations, not behavioral claims; they are held at their baseline values, and Table 14 gives their governance classes and ranges.

Exit (selection). Exit works differently for each entity, and the rules are stated in full because they decide the survival results. A firm is *distressed* when its loss streak reaches three years and its cash falls below -1.5. It shows *no traction* when its share falls below 0.01. A legacy parent exits on distress alone; the exit rule has no portfolio exemption, although dividends from acquired and alliance units enter the parent's cash. An emerging independent exits when its loss streak reaches three years and it either lacks traction or has cash below -1.5; the traction condition is what removes sub-scale firms in the shakeout, and switching the share floor off is one of the knockouts of Section 4.8. A sponsored affiliate is shielded within its route's patience window (six years of age for a build, five for a partner or acquired unit; because ages advance just before selection, the last year-end at which a unit is shielded is its fifth for a build and its fourth for the others) unless its sponsor has exited and it is distressed; after the window it exits on distress or no traction. These grace periods protect young affiliates through the liability of newness, the higher death rate of organizations early in life (Jovanovic, 1982; Stinchcombe, 1965). We avoid a single uniform exit rule because it would either kill ventures before they could ramp up or protect parents whose core business had already failed. Corporate survival is reported separately from emerging-market success.

3.6 Outcomes and links to phenomena

Every result in Section 4 is stated in one of the measures defined here, and each definition names its population and denominator, because several of the paper's distinctions turn on them.

Route and entry. A parent's *first route* is the route of its first executed entry; a rejected alliance offer or an attempt that fails for lack of cash does not set it. After a three-year cooldown a parent can enter again: the first route labels the parent in the route-share and by-route measures, while the portfolio demand behind the outcome class, and the affiliate counts, include every affiliate the parent creates. The *route share* of build, partner, or acquire is the fraction of all legacy incumbents whose first route was that route, and the *no-entry share* is the fraction that never entered; the four shares sum to one over the whole population of incumbents, entrants and non-entrants alike.

Feasibility and preference. A route's *feasibility* is, among the firm-years in which the attention gate opened, the fraction in which the route was affordable and ready by Eq. 16; its *preference* is, among those firm-years in which it was feasible, the mean choice probability it received in Eq. 17. The two separate a change in what firms can do from a change in what they want.

Outcome class. Portfolio value is θ-weighted: owned units count fully, alliance units at the equity stake $\theta = 0.5$, with demand, profit, and investment all weighted consistently. At the end of a run each entered parent is classified by a single route-blind rule: *leader* if its θ-weighted portfolio demand reaches the 0.75 quantile of entered parents and is positive, *survivor* if it reaches a floor of 0.35 demand units (total demand starts at $D_0 = 100$), else *failed*; a parent that never entered is *no entry*. Both thresholds are applied identically across routes, which is the last of the symmetry choices. Because a leader is by definition in the top quartile of entered parents, the *leader rate*, the fraction of all incumbents classified leader, is close to one quarter of the entry rate, and a difference

in it across conditions reflects entry as well as sorting among entrants. The *leader-or-survivor rate* is the fraction classified leader or survivor: over all incumbents when it describes a run, and over entered parents when it is grouped by first route in the survival experiment. *Corporate survival* is the fraction of incumbents still active at the end of the run; a parent exits on distress in its core business alone.

Affiliate persistence. For each route, the number of that route's affiliates still active at the end of the run divided by the number ever created, computed per run and averaged. It is an affiliate-level measure and ranks the routes differently from the parent-level survival rate.

Timing. The *pioneer advantage* is, within a run, the leader rate of early entrants minus the leader rate of late entrants, averaged across runs. In the baseline and under the K3 knockout the split is at the median entry year of that run's entered parents; under K5 it is the randomly assigned wave, with the late wave barred from entering before year 12.

Resources and inequality. The *resource tercile* is a descriptive label computed once at initialization for incumbents only, from the sum of the standardized values of $\log(1+F)$, K, T, L, and $\log(1+C)$, cut into thirds; route shares and leader rates by tercile use the incumbents in that tercile as the denominator. The *outcome Gini* is the Gini coefficient of θ-weighted portfolio demand across entered parents.

Market structure. Concentration is the Herfindahl index of emerging demand shares in the final year; the *shakeout drop* is the peak-to-final fall in the count of active emerging firms as a fraction of the peak.

The phase map. A phase map colors each cell of a two-parameter grid by its *winning mode*: the route with the largest mean share of incumbents that end the run as leaders having taken that route first, averaged over the replicate runs of the cell. Waiting cannot win, because it produces no leaders, and the map does not condition on entry. The map therefore says which route produces the most leaders under each condition. It says nothing directly about which route is most often chosen or which route survives best per firm, which are the separate measures above. A tie would go to the first of build, partner, and acquire, and a cell with no leaders would be left uncolored; neither occurs in the reported grids. Because a leader is a fixed top quartile of entrants, a route chosen more often tends to produce more leaders, so the map reflects choice frequency as well as success among entrants. A mode's *area share* is the fraction of a grid's cells it wins. Table 6 sets the three questions side by side.

Table 6: Three questions the results keep separate, with the exact measure behind each. They need not agree, and part of the paper's contribution is in telling them apart.

Question	Asked about	Measure	Where
Feasibility and preference	Firm-years in which the attention gate opened	Feasibility: share of those firm-years in which the route was affordable and ready. Preference: mean choice probability of the route where it was feasible	SQ1, Sec. 4.2
Regional dominance (the phase map)	A cell of the conditions plane	Winning mode: the route with the largest mean share of incumbents ending as leaders by first route, over the cell's replicate runs; summarized as area shares. The per-entrant variant divides each route's leader share by its route share (Table 17)	SQ4, Sec. 4.5
End-of-run outcome of entrants	Parents that entered, grouped by first route	Leader-or-survivor rate among entered parents, per route and condition	Sec. 4.7

Table 7 then runs from each phenomenon the paper studies to the rules that generate it, the measure in which it is observed, and the manipulation that tests whether the named rules are responsible. Three points follow from the mapping. First, the market-level regularities the model is checked against, the transition itself, rising concentration, and the shakeout, are produced by two rules, the competitive-substitution rule of Eq. 3 and the sharpened allocation of Eq. 8 with its exit floor, and Section 4.8 tests this by switching the sharpening off, removing the floor, and decoupling concentration from the dominant-design index. Second, the firm-level regularities the paper's questions are about, who searches, which route is taken, and what survives, are produced by the decision rules of Section 3.4 acting on a menu that acquisitions deplete and entry replenishes (Eqs. 18 and 9); that coupling makes the model an interacting population, and it is what the K1 knockout removes. Third, the hypothesized erosion of the first-mover advantage with pace requires the imitation and obsolescence channels inside Eq. 20, which are switched on only in the timing experiment and in the first-mover arm of the knockout panel. At baseline, with those channels off, the gap rises with pace instead (F5, Section 4.1).

Table 7: From phenomena to rules, measures, and tests. Each phenomenon the paper studies, the rules that generate it, the exact measure in which it is observed (defined above), and the manipulation that tests whether the named rules are responsible.

Phenomenon of interest	Generated by	Observed as	Identifying test
Changing opportunities for build, partner, acquire, and wait	Attention and search (Eqs. 10–11); candidate availability and acceptance (Eqs. 12–13); costs and feasibility (Eqs. 15–16); execution and depletion (Eq. 18); replenishment by entry (Eq. 9)	Route shares over all incumbents; feasibility and preference over evaluated firm-years	One-factor sweeps of pace, target liquidity, partner reliability, slack, integration friction, and uncertainty; clean K1, in which acquisitions copy the target instead of removing it while entry stays on (Sec. 4.2)

Table 7 (continued)

Phenomenon of interest	Generated by	Observed as	Identifying test
Early-entry advantage	Entry timing set by the attention gate; time in market compounding through learning (Eq. 20) and the rising barrier (Eq. 4); imitation and obsolescence channels switched on only in the timing experiment and the first-mover arm of the knockout panel	Pioneer advantage: early minus late leader rate, median split (baseline, K3) or wave split (K5), at six pace levels	K3, a coin flip replaces the attention gate; K5, random early and late waves with the late wave barred until year 12 (Sec. 4.3)
Resource sorting and unequal outcomes	Correlated initialization (Eq. 1); affordability (Eq. 16) and cost relative to cash (Eq. 17); subsequent competition under Eqs. 8 and 20	Route shares and leader rate by initial resource tercile; outcome Gini over entered parents	Tercile comparison at baseline; K4, all firms started at the attribute means, for the Gini only (Sec. 4.4)
Conditional phase-map regions	Route values (Eq. 14), costs (Eq. 15), feasibility (Eq. 16), and choice (Eq. 17); contractibility through acceptance (Eq. 13) and the internalization premium	Winning mode per cell and area shares; Sobol first-, total-, and second-order indices over eleven parameters	Four two-parameter grids at three replicate runs per cell over fifty years; Sobol at 1,024 base samples; contractibility ablation; acquisition-cost and acquisition-friction sweeps (Secs. 4.5, 4.8)
Concentration and shakeout	Competitive position (Eq. 7) and sharpened allocation (Eq. 8); exit rules (Sec. 3.5); the closing entry window (Eq. 9)	Final emerging HHI; shakeout drop	Sharpening switched off ($\gamma = 0$); share floor removed; decoupled-HHI variant (Sec. 4.8); calibration targets F1 and F2 (Sec. 4.1)
Route-dependent persistence and success	Affiliate starting state (Eq. 18); learning (Eq. 20); parent support, patience windows, and exit rules (Sec. 3.5)	Affiliate persistence by route; leader-or-survivor rate by first route among entered parents; logistic comparison net of initial tercile and seeded size class	Liquidity conditions (Sec. 4.6); seven-condition survival experiment and the build-scale sweep (Sec. 4.7)
The transition takes off, stalls, or fails	Competitive substitution on the signed frontier gap (Eq. 3); population-derived frontiers; learning that moves them (Eq. 20)	Path of m_t; disruption flag	Pace sweep; calibration target F1 (Sec. 4.1)
Alliances under uncertainty	Option value and real-options discount (Eq. 14); partner cost rising in u_t (Eq. 15)	Partner share (over all incumbents) at low and high uncertainty	Uncertainty sweep; stylized fact F3, not reproduced (Sec. 4.1)

3.7 Experiments and evidentiary interpretation

The pipeline defines one world, and the experiments find and stress-test the contingencies that govern it. Six experiments map onto the four subquestions, and a knockout panel supplies identification. Table 8 gives the matrix with the settings a reader needs to interpret each comparison.

We reuse random seeds across conditions (common random numbers) to reduce simulation noise. Because a single seeded stream feeds each run, this aligns the starting populations but only partly aligns the random draws after that: a swept parameter that changes how many random draws happen early throws the stream out of step later on, so the variance reduction is a lower bound, not a clean cancellation. We therefore report outcomes as distributions across many runs with confidence intervals, never as a single representative run.

Table 8: Experiment matrix. Every condition uses common random numbers; outcomes are distributions across runs with 95% confidence intervals. Timing channels (imitation rate 0.12, imitation damping 12, obsolescence gain 10) are switched on only in Experiment 2 and in the first-mover arm of the knockout panel.

Experiment	Question	Manipulation and levels	Runs per cell	Horizon (years)
1. Route mix	Do conditions change the most common route, and through feasibility or preference?	One-factor and factorial sweeps: pace a 1.5/4.0; target liquidity ρ_{entry} 1.5/5.0; partner reliability θ_{acc} 0.78/0.45; slack κ_F 60/15; integration friction ℓ_0 0.05/0.45; uncertainty low (stage values 0.15 at emergence and 0.10 in growth, no shocks) / high (0.45 and 0.35, shock probability 0.10)	200	60
2. Timing and pace	Is the early-entry advantage a timing effect or a selection artifact, and how does it vary with pace?	Pace $a \in \{0.8, 1.2, 1.8, 2.6, 3.6, 4.6\}$ in three arms: endogenous timing, K3, K5	200	60
3. Sensitivity and phase maps	Which parameters set the region boundaries, and which route produces the most leaders where?	Sobol over eleven parameters: pace a, frontier-gap weight b, de novo entry rate ρ_{entry}, acceptance center θ_{acc}, integration loss ℓ_0, costs c_b and c_p, price multiple μ, equity stake θ, initial emerging frontier, and contractibility χ (1,024 base samples, two replicates each); four two-parameter grids of 12×12 cells (three replicates per cell)	2 per Sobol point; 3 per grid cell	50
4. Architecture	What market composition and affiliate persistence emerge?	Target liquidity ρ_{entry} 1.5 / baseline 3.0 / 5.0	200	60
5. Sorting	Do resources sort firms into routes, and is inequality inherited or produced?	Baseline versus K4	200	60
6. End-of-run outcome	Which first route has the highest end-of-run leader-or-survivor rate, under which conditions, and does matching mode to condition help?	Seven conditions: baseline; χ 0.2/0.8; a 1.5/4.5; uncertainty at emergence 0.18/0.42	200	60

Knockouts (200 runs each, sixty years). K1 clean: an acquisition copies its target instead of removing it (the copy starts at age zero with a cash seed of half the route cost), entry stays on, at $\rho_{\text{entry}} = 1.5$ (P1). K3: each parent-year a coin flip with probability 0.5 replaces the attention gate; the three-year cooldown still applies (P2). K4: every firm, entrants included, starts at the attribute means (P3). K5: each incumbent is assigned at random to an early or late wave, the late wave barred until year 12, both waves passing the same gate (P2). Further switches used in Section 4.8: sharpening off ($\gamma = 0$), share floor removed, contractibility gains set to zero at $\chi = 0.9$, and an acquisition-price friction rising with dominant-design strength ($\xi > 0$). The panel's first-mover arm runs at $a = 3.4$ with the timing channels on. Robustness runs use their own designs: the fifty-percent perturbation draws 200 Latin-hypercube points over eleven parameters with three seeds each over 55 years; the cost-neutral and acquisition-friction checks use 150 seeds over 55 years; the decision-temperature and decoupled-HHI checks use 50-year runs; the build-scale sweep uses 80 seeds; the second-order Sobol analysis uses 128 base samples with four replicates over 40 years.

Following the generative standard, we report each generative, confound, and decomposition headline

both with and without the manipulation its type calls for, and the three kinds of manipulation carry three distinct expectations (Epstein, 1999; Davis et al., 2007). When the rule is *removed*, a generative pattern must collapse or fall sharply toward its null, the two verdicts reported in Table 11. When an *alternative cause is randomized away*, a confound-tested pattern must persist; if it collapses, the alternative cause was doing the work. When a *decomposition* switch is thrown, the pattern must persist while the switch reveals which channel carries its size. Table 11 in Section 4.8 reads each headline against the expectation of its own type. We back these manipulations up with Sobol sensitivity analysis, a variance-based method that splits each outcome's variation among the individual parameters and their interactions (Saltelli et al., 2010; ten Broeke et al., 2016). Sobol analysis is the right tool here because an agent-based model's nonlinear interactions can defeat a one-factor-at-a-time analysis. Second-order indices, estimated in a smaller separate design, show interactions of modest size for the acquire and partner shares and small ones for the build share and overall leadership.

Five layers of evidence support the model (Table 9) that together cover the forms of validation an agent-based study is expected to pass (Rand and Rust, 2011; Harrison et al., 2007). *Conceptual validation*, that the agents, mechanisms, and decision rules are theoretically justified, is carried by the theory-to-model map (Table 1) and the design–evidence map (Appendix C). *Face* and *pattern validation*, that the dynamics and the emergent outputs resemble documented transitions, are the stylized-fact scorecard (Section 4.1) and the empirical anchoring (Appendix D). *Internal* and *mechanism validation*, that results come from the model's logic rather than coding artifacts and that specific micro-rules produce the macro-patterns, are the knockout panel (Section 4.8). *Sensitivity, robustness, and extreme-condition tests* are the Sobol analysis, the fifty-percent perturbation, the specification variants, and the decision-temperature sweep toward its near-deterministic limit (Section 4.8). The inferential weight rests on the knockouts and robustness checks, not the calibration fit, which only places the baseline inside plausible bands; on the choices involved in calibrating and validating agent-based models see Werker and Brenner (2004), Brenner and Werker (2007), Fagiolo et al. (2007), and Windrum et al. (2007), and on drawing policy conclusions from simulation models, Brenner and Werker (2009).

Table 9: The validation taxonomy. Five layers of evidence, what each establishes, and where it is reported. The layers are separate; a no-data model is disciplined by the knockouts and robustness checks, not by the calibration fit.

Evidence layer	What it does, and what it establishes	Where
Calibration targets	The sharpening gain and the exit floor are tuned so that concentration (F2) and the shakeout (F1) land inside their bands; establishes a plausible starting point, not a fitted result	Sec. 4.1; App. B
Stylized-fact checks	Baseline is checked against six documented regularities (F1–F6); F1 and F2 are the calibration targets; of the four untargeted checks F4 passes, F3 and F5 are not reproduced, and F6 is not scored; establishes qualitative plausibility	Sec. 4.1
Mechanism knockouts	Each generative, confound, and decomposition headline is paired with a manipulation: the rule removed (generative, must collapse), the alternative cause randomized (confound, must persist), or the effect decomposed (decomposition, must split); establishes that each pattern is produced by the named rule	Secs. 4.2–4.8
Robustness checks	±50% perturbation, specification variants, cost-neutral and acquisition-friction runs, and a decision-temperature sweep; establishes the findings survive off the baseline	Sec. 4.8
Empirical anchoring	Model mechanisms are mapped to documented technology transitions; establishes that the primitives correspond to real cases, not only to internal logic	App. D

4 Results

Section 4.1 checks the calibrated baseline against documented regularities, Sections 4.2 to 4.5 take the four subquestions in turn, Sections 4.6 and 4.7 examine market composition and the payoff to entrants, and Section 4.8 collects the manipulations that test each headline. Three measures recur, and Table 6 fixes each: whether a route is feasible and preferred for a given firm, which route produces the most leaders in a region of a conditions plane, and how often a firm that entered by a given route ends the run as a leader or survivor. The three need not agree. Intervals are 95% intervals across runs, and differences between conditions that share seeds are paired by seed.

4.1 The baseline reproduces a transitioning market

The calibrated baseline reproduces the familiar arc of a regime transition. The emerging market grows along an S-curve, the disruption condition fires in most runs, concentration rises toward a dominant design, and within each run the number of emerging firms falls after its peak (Figure 2; the runs peak in different years, so the mean curve declines only gently). We score the baseline against six stylized facts, each with a target band fixed before the runs, across twelve seeds over sixty years (Table 10).

Two of the facts are calibration targets: we tuned the increasing-returns sharpening and the exit floor so that rising concentration (F2) and the diffusion-with-shakeout arc (F1) land in their bands. The other four are untargeted checks. One passes (F4), two are not reproduced (F3 and F5), and one is not scored (F6).

Table 10: Stylized-fact scorecard. Each fact has a target band fixed before the runs; the baseline is twelve seeds over sixty years. F1 and F2 are the calibration targets; F3 to F6 are untargeted checks; F6 is not scored.

Fact	Pattern	Target band	Baseline	Status
F1	Diffusion with a post-peak shakeout	disruption fires ≥ 0.6; share $\in [0.5, 0.95]$; firm-count drop ≥ 0.20	fired 0.75; share 0.65; drop 0.32	in band (calibration target)
F2	Rising concentration to a dominant design	emerging HHI rises; final $\in [0.03, 0.5]$	rise +0.16 to 0.23	in band (calibration target)
F3	Alliances favored under uncertainty	partner share higher under high uncertainty	0.18 (low) vs. 0.18 (high); Exp. 1: 0.184 vs. 0.182	flat; not reproduced
F4	Acquisitions for distant capability gaps	acquirer frontier gap > builder's	0.188 vs. 0.182	pass
F5	Pace-contingent first-mover advantage	advantage peaks at intermediate pace	rises with pace: 0.121, 0.141, 0.168 (slow, medium, fast)	not reproduced
F6	Fragility of shared-control units	alliance survival below owned when instability is on	not scored	n/a

Because F1 and F2 were tuned, their fit is not independent confirmation (Klepper, 1996; Anderson and Tushman, 1990). What they add is the knockout evidence of Section 4.8: concentration collapses when the sharpening is removed, and the shakeout collapses when the exit floor is removed.

F4 passes: acquirers are further behind the frontier than builders (0.188 against 0.182), in line with evidence that capability gaps push firms toward external sourcing (Capron and Mitchell, 2009). The pattern follows from the valuation rule, in which building uses the firm's own technology and acquiring uses the target's, so it confirms that the rule works as intended rather than supplying independent evidence. F3 is not reproduced: the partner share is 0.184 at low and 0.182 at high uncertainty, a comparison Section 4.2 returns to. F5 is not reproduced either. Its band expects the first-mover gap to peak at an intermediate pace, but at baseline, with the imitation and obsolescence channels off, the gap rises with pace (0.121, 0.141, and 0.168 at slow, medium, and fast pace). F6 is not scored.

The two failed checks mark real limits of the baseline: its hedging channel is weak, and without the timing channels its first-mover gap grows with pace rather than eroding. Neither fact is an input to the route-choice results, and the experiments that vary uncertainty (Experiment 1) and pace (Experiment 2) examine both in more detail (Sections 4.2 and 4.3). The baseline therefore establishes plausibility; the inferential weight rests on the counterfactual experiments, the sensitivity analysis, and the knockout panel.

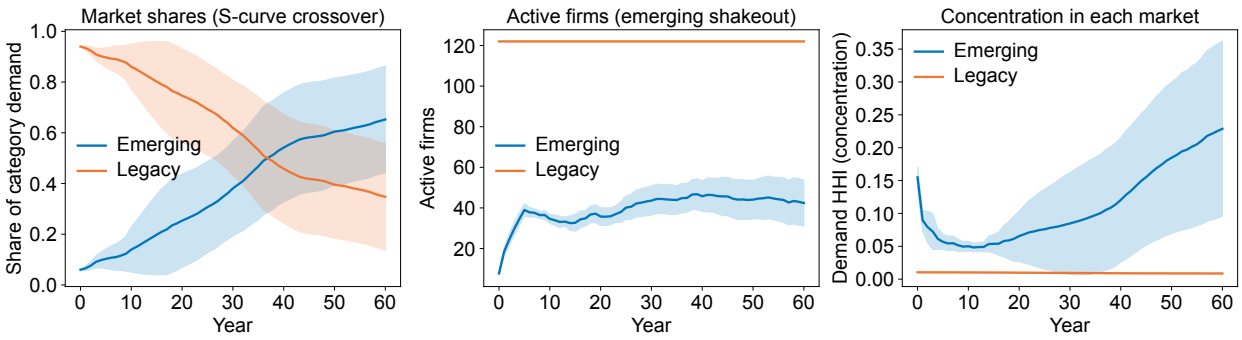


Figure 2: Baseline dynamics in both markets, with 95% confidence bands across twelve seeds over sixty years. Left: the emerging share grows along an S-curve while the legacy share declines. Center: the mean active emerging-firm count rises and then declines; because runs peak in different years, the per-run peak-to-final drop (0.32, Table 10) is larger than the mean curve shows. Right: emerging-market concentration rises toward a dominant design while the legacy market stays fragmented.

4.2 SQ1: signals govern who searches, but feasibility reshapes the route mix

Market conditions shift incumbents among the routes through feasibility rather than preference. The sensing half of SQ1 is answered by construction: the signals of legacy erosion feed the attention gate (Eq. 10), which decides who searches and when but not which route a searcher takes. The route mix is decided downstream.

Acquire is the most common first route at baseline (0.300 ± 0.013 of incumbents, against 0.250 ± 0.006 for build), and at low target liquidity build overtakes it (0.300 ± 0.009 against 0.221 ± 0.009). Feasibility and preference separate the reasons. As target liquidity rises, acquisition feasibility roughly doubles, from 0.093 to 0.204 (paired difference $+0.112 \pm 0.016$), while acquisition preference does not move (-0.001 ± 0.001); partner feasibility rises by $+0.160 \pm 0.018$ while partner preference moves by -0.004 ± 0.001. Firms do not change what they want; the supply of targets changes what is within reach. This is the pattern proposition P1 expects.

Real-options logic would predict a hedge instead: firms partner more when uncertainty is high (Folta, 1998). In the model they do not. Partner share is flat across the uncertainty levels of Experiment 1 (0.184 under low and 0.182 under high uncertainty; paired difference -0.002 ± 0.012). The flat response has a mechanical source: uncertainty raises the partner route's option value in Eq. 14 but also its cost in Eq. 15, and at baseline the two roughly offset. The model's hedging channel is weak, and the route mix responds to the supply of targets, not to uncertainty.

The K1 knockout examines the role of target depletion. In K1, an acquisition copies its target instead of removing it from the pool, while entry stays on, so the test separates depletion from the size of the pool. At a low entry rate ($\rho_{\text{entry}} = 1.5$), build share falls from 0.300 ± 0.009 to 0.135 ± 0.004, a fall of fifty-five percent, because acquisition feasibility rises from 0.09 to 0.40 and the number of active independents from 0.5 to 4.0. The alternative explanation, that a frozen pool starves the acquire route by leaving fewer firms, does not fit: the knockout run has *more* independents, and build still falls. The copy is not a perfect counterfactual, since it starts at age zero and leaves the original independent active, so K1 also adds capacity and a younger acquired unit.

For SQ1, signals decide who searches and when, and feasibility, set by rivals' earlier choices, decides what searchers can do. The electric-vehicle transition offers a loose parallel: after GM acquired Cruise in 2016 and Ford invested in Argo in 2017 (General Motors, 2016; Ford Motor Company, 2017), fewer autonomy startups were available to later movers (Appendix D, an illustration, not a test).

4.3 SQ2: early entry pays at every pace tested

Early entrants, those entering at or before the median entry year of their run, have a leader rate $+0.136 \pm 0.010$ higher than late entrants, averaged over 200 runs at each of six pace levels in the timing experiment; we call this the endogenous estimate. The gap could reflect timing or selection, if the strongest or most pressured firms happen to enter first. The K3 knockout replaces the attention gate with a coin flip (probability 0.5) each year, so that entry timing is cut loose from the gate that would bring the most pressured firms in first. The gap, measured with the same median split, is then $+0.145 \pm 0.011$, essentially unchanged.

K3 randomizes the entry gate but not the resources a firm brings, so stronger firms still clear feasibility and still tend to enter earlier. The K5 knockout cuts entry order loose from firm strength directly: each incumbent is assigned at random to an early or a late entry wave, and the late wave is barred until year 12, close to the typical median entry year of the endogenous runs. The two cohorts are matched on strength and differ only in when they may enter. The gap between the waves is $+0.105 \pm 0.011$, about three-quarters of the endogenous estimate and positive at every one of the six pace levels. Severing entry order from firm strength therefore leaves a time-in-market effect, although the early wave also has the target pool to itself until year 12, so part of that effect may be first access to targets rather than learning. These results support the first half of proposition P2.

The second half of P2 expects the advantage to erode as the pace rises, because a fast-moving frontier makes an early lead obsolete (Suarez and Lanzolla, 2007; Lieberman and Montgomery, 1998). Pace here is the substitution rate a of Eq. 3; the obsolescence channel reads the realized frontier progress g_t (Section 3.5), which faster substitution raises by moving demand, cash, and R&D toward emerging firms. With the imitation and obsolescence channels on, the endogenous gap stays between 0.136 and 0.148 up to $a = 2.6$ and falls to 0.117 and 0.126 at $a = 3.6$ and 4.6; the paired drops from the slowest pace, 0.031 ± 0.018 and 0.022 ± 0.018, exceed their sampling error. In the K3 and K5 arms the gap stays near 0.145 and 0.105 at every pace (Figure 3).

The erosion prediction therefore receives partial support, with two limits on its scope. The erosion appears only when firms time their own entry, which suggests that a fast pace changes which firms enter early more than it reduces the value of time in the market. And it appears only with channels that are switched on for this experiment; at baseline, with those channels off, the gap rises with pace instead (F5 in Table 10). The pattern fits an empirical record in which first-mover effects are small and depend on conditions (Lieberman and Montgomery, 1988, 1998; Suarez and Lanzolla, 2007).

The electric-vehicle transition shows a similarly mixed record (Appendix D): Tesla's lead in battery-electric vehicles has largely held, while incumbents' autonomy investments on the faster-moving frontier faltered.

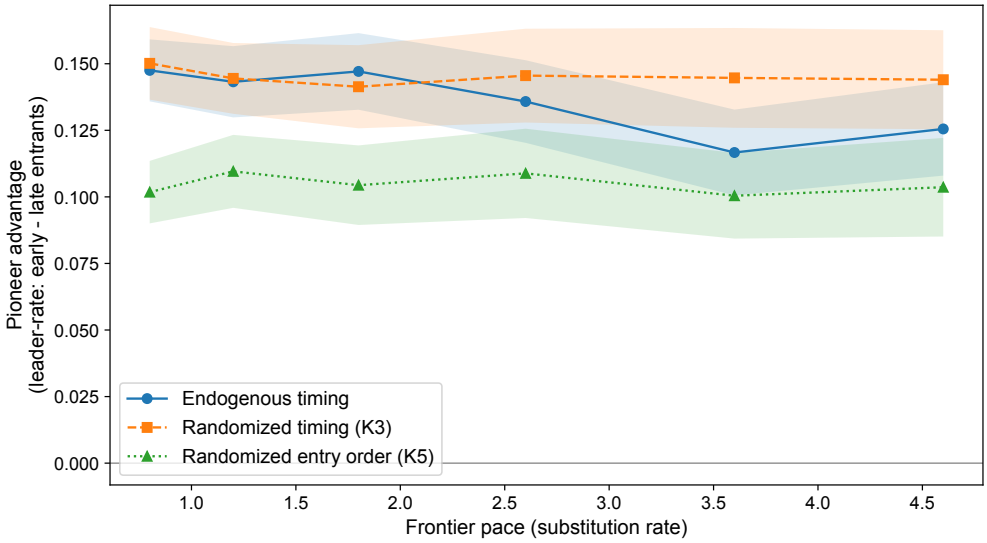


Figure 3: Pioneer advantage (early minus late leader rate) as a function of the substitution rate a, in the endogenous arm and under the randomized-timing (K3) and randomized-entry-order (K5) knockouts; lines are means and shaded bands 95% confidence intervals across 200 runs per point. The gap survives both randomizing who enters when and severing entry order from firm strength (K5 gap +0.105, positive at every pace). In the endogenous arm the gap narrows at the two fastest paces; the K3 and K5 gaps stay flat. The imitation and obsolescence channels are switched on for this experiment.

4.4 SQ3: resources sort firms, but stratification is path-dependent

Resources do sort firms into routes. Comparing the top and bottom terciles of incumbents on the initial resource index of Section 3.6, with the incumbents in each tercile as the denominator, the share whose first route was build is higher in the top tercile by $+0.190 \pm 0.013$, the partner share by $+0.057 \pm 0.011$, and the acquire share by $+0.047 \pm 0.015$; the complement is that the bottom tercile more often does not enter at all, since the four shares sum to one within a tercile. This runs against the expectation of Section 2.2 that resource-poor firms would lean toward the flexible routes: they partner less, not more, and mainly stay out. Varying slack directly shows the same pattern. Lowering it, by raising the cash normalizer κ_F from 15 to 60 in Experiment 1, raises the share of incumbents that never enter from 0.25 to 0.36 (paired difference 0.11 ± 0.02). The sorting carries into outcomes: the leader rate is higher in the top tercile by $+0.212 \pm 0.014$. This matches the capability premium that dynamic capability theory predicts and that meta-analytic evidence supports (Teece, 2007; Cohen and Levinthal, 1990; Bitencourt et al., 2020).

Outcome inequality is not only the starting endowment under another name. The K4 knockout starts every firm at the attribute means and asks whether inequality in outcomes persists. It does, and it grows: the outcome Gini rises from 0.823 to 0.864 (paired difference $+0.041 \pm 0.016$). Because the Gini is computed over entered parents, K4 also changes who is in that population, and the remaining inequality comes from stochastic draws amplified by the sharpened allocation. The comparison shows that outcome inequality persists after initial resources are homogenized, as P3 expects; it does not by itself identify the sorting channel, which rests on the tercile comparison above. For SQ3, resources shape which routes a firm takes and how often it leads, while the size of the inequality in outcomes builds up through competition. Figure 4 shows route shares and leader

rates by initial resource tercile; Table 11 reports the K4 comparison.

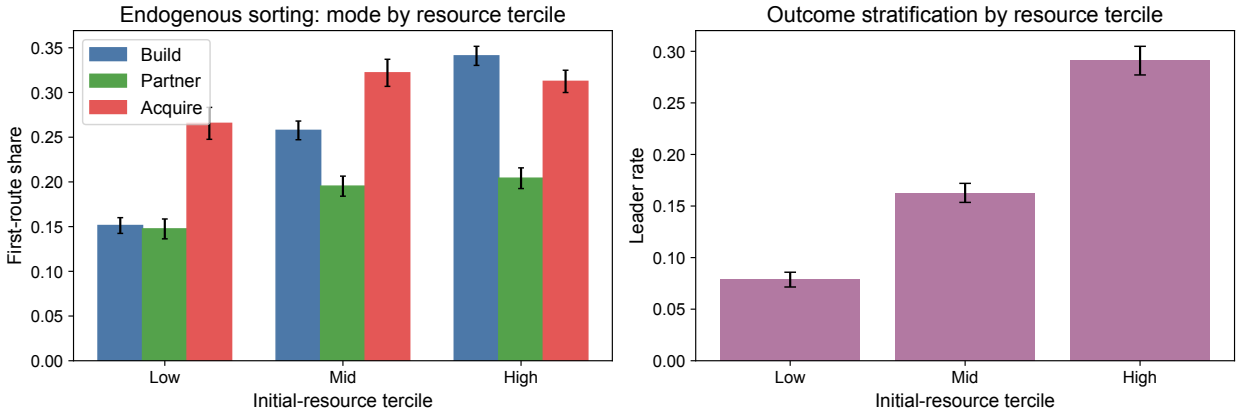


Figure 4: Sorting and stratification by initial resource tercile (baseline, 200 runs; error bars are 95% confidence intervals across runs; the denominator in each tercile is the incumbents in that tercile). Left: first-route shares by tercile; the top tercile takes build first more often by +0.190. Right: leader rate by tercile; the top tercile leads more often by +0.212. The K4 comparison, in which the outcome Gini rises from 0.823 to 0.864 when every firm starts at the attribute means, is reported in the text and in Table 11, not in this figure.

4.5 SQ4: route costs set the region boundaries

The questions so far concern which firms search, when they enter, and how they sort. SQ4 asks what shapes the mode-dominant regions. In the sensitivity analysis, the build and partner shares depend mainly on their own costs, while the acquire share depends mainly on the initial emerging frontier because the acquisition fixed cost is held at baseline. A variance-based Sobol analysis over the eleven treatment parameters traces the build share mainly to the build fixed cost, with a total-order index of 0.95, where the total-order index is the share of an outcome's variance due to a parameter together with all its interactions (Saltelli et al., 2010). The partner share traces the same way to the partnership fixed cost, with an index of 0.63. The acquire share traces mainly to the starting emerging frontier, with an index of 0.37.

Two cautions keep these indices from being read as more than they are. First, the fact that a route's share falls as its own cost rises is partly built into the cost term in the choice rule, so the indices are best read as a check that the levers move in the right direction. Second, the build index is the highest partly because the build fixed cost is swept over the widest range of the eleven parameters in raw units, so a high total-order index there partly reflects that width. The non-cost levers act on route shares more than on the winning mode: contractibility moves the partner share even with costs held equal, and target depletion moves the route mix through feasibility (Section 4.2), but at the ranges we map neither redraws the winning regions much (Figure 5). Read with these cautions, the indices point to the routes' own costs, levers of the seizing stage, as the main influences on the build and partner shares; sensing and transforming levers matter for who enters and who survives, but they do not by themselves set the boundaries between the regions.

The second-order indices, estimated in a smaller separate design (128 base samples, four replicates, forty years), fill in the picture, though their precision should be read modestly. For the build share no pairwise interaction exceeds about 0.07 (the largest pairs the entry rate with the build

cost), so the build region's sharp boundary reflects mainly one lever. For the partner share the largest interactions, about 0.09, pair the acquisition price multiple with contractibility and with the partner cost; for the acquire share the build cost interacts with the starting frontier, the price multiple, and the partner cost, each at about 0.12 to 0.13. For overall leadership the second-order indices are small and some are negative, which at this sample size signals estimation noise. The acquire share is the most spread out of the three, shared across the starting frontier, the supply of targets, and the cost of acquisition rather than governed by a single lever, which fits with its region giving way once the cost of acquisition rises above its baseline. The bootstrap confidence intervals on all of these indices are wide at the sampled resolution, so we read them as an ordering and an order of magnitude, not as precise values.

Contractibility moves the choice between partnering and owning. As contractibility rises, the partner share climbs from 0.07 to 0.27, because a more contractible capability is one partners will trade. This is not cost in disguise: with the build and partner costs held equal, partner share still rises by +0.13 across the contractibility range, so the lever works through acceptance and the ownership premium rather than through route cost. Switching the mechanism off flattens the gradient to nearly zero. Even so, on the contractibility-by-build-cost plane partnering wins only 0.03 of the cells, because acquisition at its baseline cost produces more leaders over most of that plane. The electric-vehicle stack offers an illustration (Appendix D): the contractible battery layer was sourced largely through cell joint ventures and supply contracts, while the tacit autonomy layer was acquired or built.

The phase maps put these levers together. Each cell of a 12×12 grid is colored by its winning mode, the route that produces the most leaders in that cell by the criterion of Section 3.6, averaged over three replicate runs per cell on a fifty-year horizon. On the plane of build cost against partner cost, the winning mode splits, as a share of the plane's cells, into acquire 0.66, build 0.18, and partner 0.16 (bootstrap 95% intervals over the replicate runs: 0.58–0.69, 0.16–0.22, and 0.13–0.22). On the plane of build cost against pace it splits into acquire 0.82 and build 0.18; on the plane of contractibility against build cost into acquire 0.78, build 0.19, and partner 0.03; and on the plane of target liquidity against uncertainty acquisition wins every cell. These planes hold the cost of acquisition at its baseline; Section 4.8 sweeps it as a third axis.

No mode wins the whole space in three of the four planes. In the build-cost by partner-cost plane each route owns a substantial region; in the contractibility-by-build-cost plane partnering holds only 0.03 of the cells; and in the liquidity-by-uncertainty plane acquisition at its baseline cost wins everywhere. What the maps show is therefore conditional on the cost of acquisition, and Section 4.8 shows the acquire region giving way to partnering as that cost rises. The split of the cost plane supports P4 in its cost form and answers the boundary half of SQ4: the boundaries between regions are set at the seizing stage by the routes' own costs, while contractibility and target supply move route shares but, at the ranges mapped, rarely the winning mode. Which route's entrants most often end as leaders or survivors is a separate measure, taken up in Section 4.7.

Because a leader is a fixed top quartile of entrants, the map mixes how often a route is chosen with how well its entrants do. Scoring each route instead by its leaders per entrant, in a rerun of the same grids with the same seeds (which reproduces the stored maps exactly), gives a different picture on the build-cost by partner-cost plane: acquire 0.81, partner 0.16, and build 0.03 of the cells (bootstrap 95% intervals over the three seeds 0.69–0.81, 0.16–0.25, and 0.01–0.08; Table 17 in Appendix E gives all four planes). Partnering keeps its region where it is cheap, while build's region is largely a frequency effect: where building is cheap many firms build, so many leaders are builders, but each builder is less likely to lead than each acquirer. The ordering is consistent

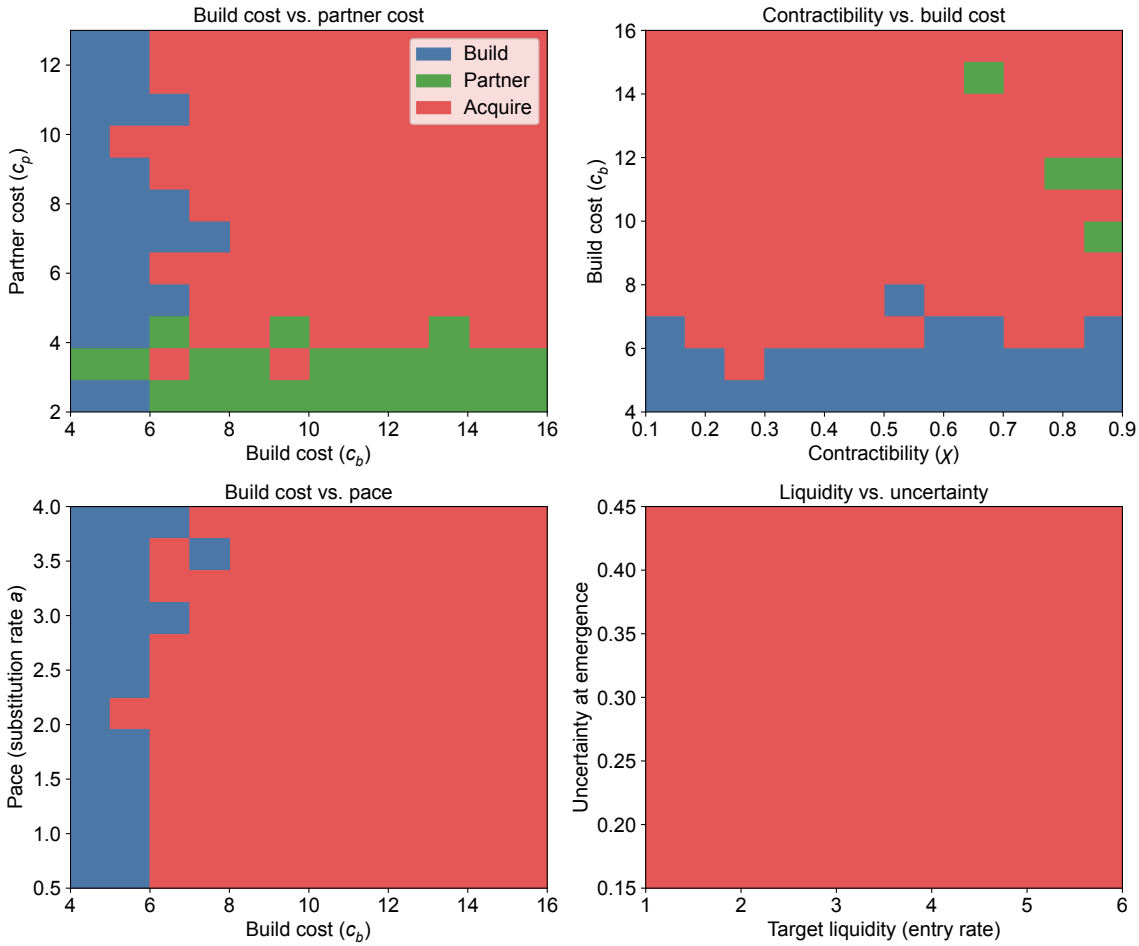


Figure 5: Phase diagrams of the winning sourcing mode, with the cost of acquisition held at its baseline. A cell's winning mode is the route producing the most leaders there (mean over three replicate runs, fifty-year horizon, 12×12 cells per plane). Area shares: build cost by partner cost, acquire 0.66, build 0.18, partner 0.16; contractibility by build cost, acquire 0.78, build 0.19, partner 0.03; build cost by pace, acquire 0.82, build 0.18; target liquidity by uncertainty, acquire 1.00. Raising the cost of acquisition redistributes the acquire region to partnering (Section 4.8).

with the leader-or-survivor ranking among entering parents in Section 4.7, although that ranking also counts survivors, whereas this map counts leaders alone. Per entrant, the ready position an acquisition delivers pays off over most of the plane, and the contingency the map shows lies mainly in which routes firms take. Figure 5 reports the maps, and Figure 8 in Appendix E reports the Sobol indices.

4.6 The emergent architecture shifts with conditions

Affiliate persistence, the share of a route's affiliates still active at the end of the run, differs sharply by route: 0.07 ± 0.01 for builds, 0.24 ± 0.01 for acquired units, and 0.38 ± 0.02 for alliances, so an alliance unit is about five times as likely as a build to be active at the end of the run.

This affiliate-level measure ranks the routes in the opposite order from the parent-level leader-or-

survivor rate of Section 4.7: alliance affiliates last longest, while parents that entered by acquisition most often end as leaders or survivors. A long-lived alliance unit and a winning corporate portfolio are different things. The composition of the emerging market also shifts with conditions: at baseline the average run ends with about 20.0 alliance units, 16.0 acquired units, and 4.3 builds, and at high target liquidity with 22.9, 15.2, and 3.3. Evolutionary work on mergers has examined how acquisitions reshape the firm-size distribution, finding a shift toward larger sizes among the firms involved but little change across the population (Cefis et al., 2009); in the model, too, every acquisition acts on the firms involved, removing an independent and enlarging an incumbent's portfolio. Figure 9, in Appendix E, reports the composition by condition.

4.7 What pays off: parents entering by acquisition end best at the calibrated build scale

Beyond which routes firms take and which route produces the most leaders, a separate question is how often a firm that entered by a given route ends the run as a leader or survivor. At the calibrated build scale, the ranking is the same in all seven market conditions (the baseline, and contractibility, pace, and uncertainty at emergence each at two levels): acquisition is highest, between 0.32 and 0.39; partnership is in the middle, between 0.29 and 0.33; and building is lowest, between 0.15 and 0.20. Each rate is pooled over the 200 seeds of its condition, and its 95% interval is at most ± 0.02.

The ordering is not only a wealth effect. A logistic model of the leader-or-survivor outcome on first route, condition, initial resource tercile, and seeded size class (both measured at initialization, Section 3.6) keeps the same ordering: relative to acquisition, the estimated log-odds are -0.23 for partnering and -1.16 for building (point estimates from a lightly penalized fit).

The source of the ordering is the scale at which each route starts. An acquisition enters at the target's full operating scale, while an internal build enters at a fraction of the parent's scale and so falls below the exit floor more often while it ramps up. Raising the build scale should then move the top of the ranking to another route.

Sweeping the build scale, the fraction of the parent's capacity at which a build starts, confirms this. Acquisition is highest at the calibrated scale of 0.22, with a leader-or-survivor rate of 0.376 against partnership's 0.307 (the sweep uses its own runs, so these rates differ slightly from the baseline row of Figure 6); the two are level at 0.340 at a scale of 0.33; and partnership is ahead at 0.359 against 0.324 at a scale of 0.40. No route has the highest leader-or-survivor rate at every build scale, and a route can produce the most leaders in a region of the map without having the highest rate among its entrants.

The leader-or-survivor ranking should also be read against the empirical record on acquisitions. That record finds that acquirers on average do not gain, and often lose, value, with failed integration a recurring cause (King et al., 2004; Haspeslagh and Jemison, 1991); the most recent synthesis shows that performance depends on the deal and its context rather than on acquisition as such (King et al., 2021). The model does not contradict that record. Its acquisition advantage already depends on the integration loss the model charges and on the absorptive capacity needed to make an acquired position pay off. The result is narrow: at the calibrated integration friction, acquiring an established, running business is the most reliable way to clear the exit floor. It is not a claim that acquisitions create value on average.

A final test concerns matching. We tested the classic prediction that firms which match their route to the condition that favors it more often end the run as a leader or survivor, counting a

firm as matched if it built or acquired when contractibility was low, partnered when it was high, built when the pace was slow, or acquired or partnered when it was fast. The matched-versus-mismatched difference is indistinguishable from zero, a premium of −0.01 (95% interval ±0.01). The contingency the maps show lies in which routes firms take and how many leaders each produces, while the end-of-run outcome of an entrant favors the ready position an acquisition delivers across the market conditions tested and at the calibrated build scale. Figure 6 reports the rates by route and condition.

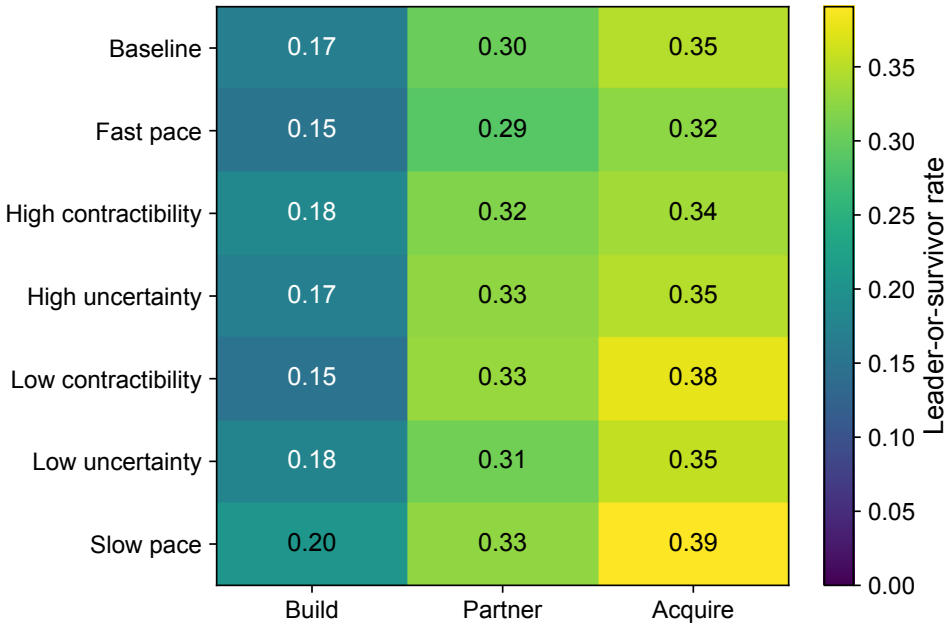


Figure 6: Leader-or-survivor rate by first route and condition. Each cell is the leader-or-survivor rate among entered parents whose first route was the column's route, pooled over the 200 seeds of the row's condition, at the calibrated internal-build scale. Acquisition has the highest rate and building the lowest in every row. The logistic comparison net of the initial resource tercile and seeded size class, the build-scale sweep in which partnering overtakes acquisition, and the null matching premium are reported in the text, not displayed here.

4.8 Identification: the knockout panel

Each generative, confound, and decomposition headline is reported with and without the mechanism that should produce it and held to the expectation for its type (Epstein, 1999); the entrant ranking is examined instead by the build-scale sweep of Section 4.7. Table 11 summarizes the core panel; stress tests and objections follow.

Table 11: Identification summary. Each headline in the table is paired with the manipulation of its candidate mechanism. Generative patterns must collapse or fall sharply when the mechanism is removed; the confound pattern must persist when its alternative cause is randomized away; the decomposition pattern must persist while revealing its channel. Supporting checks and stress tests follow in the text.

Headline	Type	Manipulation	Key result	Verdict
Rising concentration	Generative	Increasing-returns sharpening off	HHI 0.240 ($\pm$0.038) $\rightarrow$ 0.017	collapses
Post-peak shakeout	Generative	Exit floor removed	Firm-count drop 0.299 ($\pm$0.040) $\rightarrow$ 0.000	collapses
Depletion-driven build	Generative	Target depletion disabled (clean K1, $\rho_{\text{entry}} = 1.5$)	Build share 0.300 ($\pm$0.009) $\rightarrow$ 0.135 ($\pm$0.004)	falls 55%
Contractibility split	Generative	Contractibility lever ablated	Partner share at $\chi = 0.9$: 0.270 ($\pm$0.013) $\rightarrow$ 0.182 ($\pm$0.009); gradient ≈ 0	falls 33%; gradient flat
First-mover advantage	Confound	K3: timing randomized; K5: entry order randomized	Gap +0.136: K3 +0.145, K5 +0.105 (Exp. 2, means over six paces)	persists
Outcome inequality	Decomposition	K4: resources homogenized	Gini 0.823 $\rightarrow$ 0.864	persists[a]

Baseline values are means over the panel's 200 runs, except the depletion row (at $\rho_{\text{entry}} = 1.5$) and the first-mover row (Experiment 2); the scorecard of Table 10 uses twelve seeds, so its values differ slightly.

[a] Outcome inequality persists, and grows, after initial resources are homogenized; the comparison does not by itself identify the sorting channel.

The generative patterns collapse or fall sharply. Of the four generative headlines, two collapse and two fall sharply when their mechanism is removed. Concentration collapses when the increasing-returns sharpening is switched off, with the Herfindahl index falling from 0.240 to 0.017. The shakeout vanishes when the exit floor is removed, with the firm-count drop falling from 0.299 to 0.000. The two are linked: in a separate sweep of the sharpening gain, the shakeout drop rises with the gain, from 0.065 at the weakest gain tested ($\gamma = 1$) to 0.35 at the strongest ($\gamma = 12$), so the sharpening drives both concentration and the shakeout, and the exit floor turns falling share into exit. The depletion-driven build and the contractibility split complete the set: build share falls fifty-five percent under the depletion knockout, and partner share falls thirty-three percent when contractibility is switched off.

The concentration metric is not feeding itself. Because the Herfindahl index also feeds the dominant-design term that drives the sharpening, we check that the concentration result is not just a mechanical side effect of that link. A decoupled knockout removes the index from the term and puts its weight on the share and the frontier instead. Concentration still arises, with the Herfindahl index reaching a median of 0.54, well above the 0.017 of the sharpening-off knockout, above the baseline itself, and above the upper end of the F2 band, so the calibrated F2 fit holds only with the index inside S_t. The link therefore dampens rather than inflates the concentration we report: the index starts low and holds the dominant-design term down early, so the rise comes from the mechanism. How much concentration appears varies across runs, which is why concentration enters validation as a band.

The confound and decomposition headlines behave as required. The first-mover advantage persists under both the K3 timing randomization and the K5 randomized-entry-order test (Section 4.3). Outcome inequality persists, and grows, when initial resources are homogenized under

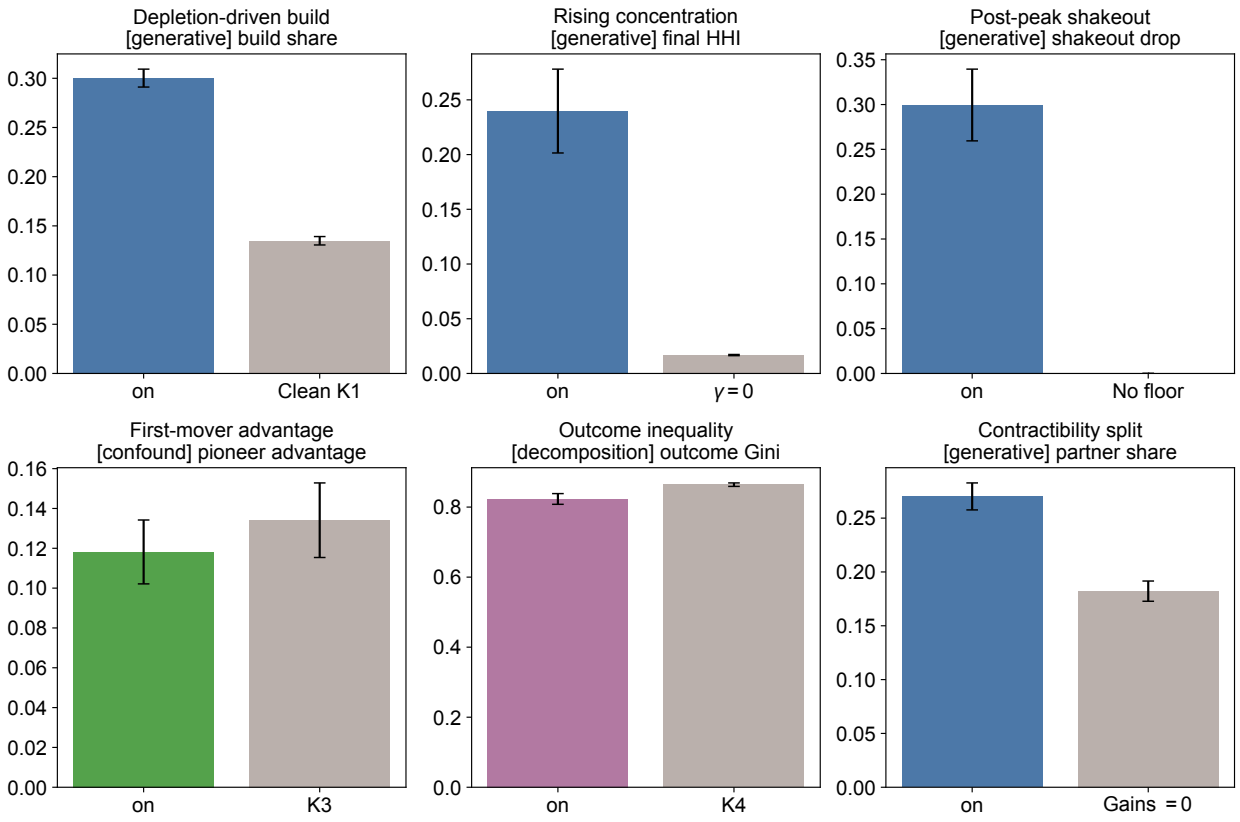


Figure 7: Knockout panel (200 runs per arm; error bars are 95% confidence intervals across runs). Each headline is shown with its mechanism on and under the named intervention; bar color marks the test type (blue generative, green confound, purple decomposition), not a route. Two generative patterns collapse and two fall sharply when their rule is removed; the confound headline (the timing experiment's channels on, at the baseline pace $a = 3.4$) persists when timing is randomized by K3, with the K5 comparison in Figure 3; outcome inequality persists under K4.

K4, which follows from the increasing-returns mechanism. Figure 7 assembles the panel.

The verdicts survive parameter perturbation. A pattern can pass its own knockout and still be an accident of the exact parameter values we used, so we push each verdict off its baseline. Under a global fifty-percent perturbation, 200 Latin-hypercube draws over eleven parameters (the pace a, the initial emerging and legacy frontiers, base obsolescence ω_0, the sharpening gain γ, the scale-economy exponent, the exit floor, the entry rate ρ_{entry}, the irreversibility penalty, the partner option-value weight, and the alliance coordination penalty), a headline survives a draw if its pattern clears a fixed threshold (for example, final concentration above 0.08 or a shakeout drop above 0.10). Resource sorting and route diversity survive in every draw, while the shakeout, the first-mover advantage, and concentration survive in 0.64, 0.71, and 0.53 of draws. For two of these the lower survival has an identifiable source: the shakeout and concentration both rise monotonically with the sharpening gain in a separate sweep, so draws that push the gain down weaken them, as they should. The first-mover advantage in this test is measured with the timing channels off.

Specification variants and the bounded-rationality premise. The five headlines of the perturbation test, which concern route shares and market structure, not the regions of the phase

map, hold under each of four specification variants: two alternative leadership thresholds (the 0.85 and 0.95 quantiles) and a halved and a doubled decision temperature ($\tau = 0.165$ and 0.66). Each mode keeps its corner of the phase map for alliance coordination penalties from 0 to 0.30; at 0.45, the top of the governed range, the partner corner passes to acquisition.

The contingency map is also stable across the governed range of the decision temperature. Sweeping it in fine steps from 0.20 to 0.50, at least two modes own regions at every level, and all three own regions across the middle of the range, so no single mode wins everywhere over the whole plausible span of bounded rationality. Pushed toward the best-response limit below that range ($\tau = 0.05$), where almost every firm takes its single highest-scoring route, the multi-region structure narrows toward one mode. Firms remain heterogeneous in those runs, so heterogeneity alone does not sustain the regions; the contingency depends on choice noise, and we list this among the limitations in Section 6. That strategic diversity follows from bounded rationality fits the behavioral and evolutionary theory of the firm (Cyert and March, 1963; Nelson and Winter, 1982).

The seeded size structure is not load-bearing. The starting size mix is held to the same standard. The contingency split, the resource-sorting gradient, and the leader-or-survivor ranking all persist as we sweep the gap between the seeded size classes from its baseline down to zero and up by half, and under alternative size mixes, so none of the headlines is an artifact of the particular mix. Two further checks confirm that the seeded class is not destiny: the class alone predicts a firm's final outcome no better than always guessing the most common outcome, with a pseudo-R^2 near 0.01, and the leader-or-survivor ranking is unchanged whether or not the size class enters as a control. Size acts through the continuous resources it sets.

The route mix is not a cost artifact. Two tests answer the most natural objection: that the route mix might just be an artifact of the route costs. With the build and partner costs held equal, the mix still moves: partner share rises by +0.133 across contractibility, a lever that touches no cost, and build share falls by 0.156, from 0.490 to 0.334, as target liquidity rises, which works through depletion. With costs held equal, contractibility and target supply still move the route mix, so the mix is not just cost by another name, although, as Section 4.5 shows, they move route shares more than they move the winning regions.

Acquisition's prominence is conditional. A second test raises the price of acquisition as the dominant design consolidates, through the factor $1 + \xi S_t$ of Eq. 15, a stand-in for merger review and target scarcity. Acquisition's share then falls from 0.297 to 0.073 while partner rises to 0.603. Sweeping the cost of acquisition across the whole choice plane confirms the point for the map as a whole. As acquisition is priced from its baseline up through a higher range and the friction is added, the share of the build-cost by partner-cost plane that acquisition wins (in this sweep's own runs, so the starting shares differ slightly from those of Figure 5) falls from 0.67 to 0.31, then to 0.08, and finally to zero, while partnership rises from 0.12 to 0.80 and building holds its own low-build-cost corner near 0.20 throughout. The acquire-heavy baseline map is the picture under cheap, unconstrained acquisition. The termination of Adobe's planned acquisition of Figma in December 2023, for lack of a clear path to regulatory approval (Appendix D), is a boundary case of this constraint, the review friction the model adds as a rising acquisition cost. Figure 10, in Appendix E, reports both tests.

Choice and survival are distinct objects. Finally, the choice map and the leader-or-survivor ranking are not one result counted twice. Across the build-cost by partner-cost plane, the route that produces the most leaders and the route with the highest leader-or-survivor rate agree in about a third of cells, and the rank correlation between the two measures is 0.09, so they carry different information. The leaders-per-entrant map of Section 4.5 sits between them: it agrees

with the most-leaders map in 0.56 of the cells of the build-cost by partner-cost plane and, like the leader-or-survivor ranking, favors acquisition over most of the plane. Taken together, the panel shows that each generative, confound, and decomposition headline behaves as its type requires; the leader-or-survivor ranking has no knockout of its own, and the test we report for it is its dependence on the build scale (Section 4.7).

5 Discussion and contributions

Dynamic capabilities are hard to observe because the construct lives in the links between processes. A generative model makes those links concrete: sensing decides who searches, seizing decides which route is attempted, and transforming decides which affiliates survive, with each stage feeding the conditions the next firm reads. Run forward, these links produce a market in which no route produces the most leaders under every condition.

5.1 Theoretical contributions

The results support three points, which together make up the dynamic account of capability sourcing the paper contributes.

First, rivals' decisions change the sourcing opportunities open to a firm. The build–borrow–buy framework treats the route choice as a firm's own decision given its resources. In the model the menu is shared. Every acquisition removes a target from the pool that later movers draw on, and the route mix responds to that depletion through feasibility rather than taste: as target liquidity rises, acquisition feasibility roughly doubles while its preference does not move, and when depletion is switched off the build share falls by more than half. Sensing therefore does not translate straight into a response. Its effect on seizing runs through the feasibility of the routes, which rivals' earlier choices have already set. This is the sense in which capability sourcing is an ecological problem, not a firm-level optimization.

Second, timing and the interactions that follow a commitment shape what the commitment is worth. In the model an early entrant compounds its position through learning and through the barrier its presence raises, and the resulting leadership gap survives both randomized timing and randomized entry order (Section 4.3). Whether the gap erodes with pace depends on who chooses the timing: with the imitation and obsolescence channels on, it narrows at the fastest paces when firms time their own entry, but not when timing or order is randomized. A pioneer's advantage can reflect time in the market or the kind of firm that moves first, and the model suggests that fast frontiers act mainly on the second. Seizing is thus a decision about when as well as which.

Third, the route most often chosen, the affiliate that lasts longest, and the parent that ends best need not share a ranking. Acquisition is the most common first route at baseline and produces the most leaders over most of each map; alliance affiliates persist longest (Section 4.6); and among parents that entered, acquisition most often ends as a leader or survivor at the calibrated build scale. Scored per entrant rather than by the number of leaders, the map itself tilts further toward acquisition (0.81 of the build-cost by partner-cost plane), and build's region turns out to reflect how often firms build. These differences matter for how sourcing is studied. In the model, affiliate persistence favors alliances, leader counts reflect both route frequency and entrants' likelihood of leadership, and leader-or-survivor rates among entering parents favor acquisition at the calibrated build scale. Each answers a different question, and the resource-pathways question,

which route fits a given firm, is closest to the comparison among entrants.

The account is also operationalized: sensing, seizing, and transforming are written as rules that can be run, so mechanisms that are hard to observe, such as how sensing turns into feasible options, can be examined one at a time. The phase map adds a comparison of modeled conditions and outcomes, a statement of which route produces the most leaders under which combination of route costs. Together these extend dynamic capability theory from firm-level adaptation to the way competition shapes how a market evolves.

The framework could be adapted to settings with more than one legacy market or emerging domain, and to questions beyond the firm. The acquisition-friction experiment shows the kind of question it can address: raising the price of acquisition as the design consolidates pushes incumbents from acquiring toward partnering, and not out of entry. The same framework can be read by competition regulators weighing a merger-review counterfactual where post-merger data do not yet exist, by platform and ecosystem owners (Jacobides et al., 2018) asking which capabilities will be partnered for rather than built, and by innovation-policy designers asking how an incentive would shift the route mix.

5.2 Methodological contribution: generative foresight with knockout identification

The study also contributes a mechanism-based approach to prediction. Instead of estimating correlations from historical data, the model makes predictions by specifying how firms decide and then simulating how those decisions play out. This approach is useful where empirical methods are weak: in a forming regime where past data are a poor guide. In place of point forecasts, the model compares route feasibility, leader production, and end-of-run outcomes among entering parents across modeled conditions.

The distinguishing move is knockout identification. A mechanism-based prediction is only as strong as the means of telling a pattern the model generates from one set by its parameters. The study removes the rule behind each generative pattern, randomizes the alternative cause behind each confound, and decomposes each effect whose channel is in question (Section 4.8). Because those headlines behave as their types require, and those tested under a wide perturbation of the parameters survive in 53 to 100 percent of draws, each pattern is tied to the mechanism named for it by the test its type calls for. In this form, generative foresight disciplines each claim with a manipulation. The contribution does not replace empirical analysis; it complements it, giving a transparent, systematic way to make and test predictions in settings where the data to estimate from do not yet exist.

5.3 Managerial implications: conditions, not prescriptions

Responding to disruption forces managers to make linked decisions: when to act, how to respond, and whether the response can be sustained. Our findings offer four practical points on these decisions, each stated as a condition to read.

First, time in the market carries value in the model, and the pace of the transition changes how much. The early-entry gap survives randomized timing and randomized entry order and stays positive at every pace tested, although it narrows at the fastest paces when firms time their own entry (Section 4.3). This is a measured comparison between early and late entrants. It is not a

general case against waiting, which the model treats as a choice every year whose value depends on the firm's cash and the uncertainty it faces.

Second, read the supply of targets, because rivals' acquisitions are draining it. The menu of feasible routes is not fixed. When the pool of acquirable targets thins, the most common route flips from acquire to build, and the flip is driven by feasibility, not by a change in taste. In the model, target availability behaves like a perishable resource that earlier movers use up.

Third, read your own slack. Resource-rich firms enter more often and lead more often, and low slack mainly keeps firms out (Section 4.4). For a cash-constrained firm the comparison argues against copying a route a richer rival can afford; because the spread in outcomes is path-dependent, the relevant question is which feasible commitment keeps learning going without exhausting slack.

Fourth, no route produces the most leaders everywhere, and the routes are not equally safe once a firm has crossed. In the build-cost by partner-cost plane, acquire produces the most leaders where building and partnering are both costly, build where building is cheap, and partner where partnering is cheap relative to building. Two cautions apply. Matching the route to the condition that favors it produced no leader-or-survivor premium in the model, so the map compares conditions and outcomes; it does not promise that the matched route ends better. And build's region reflects how often firms build more than how well builders do. Among parents that have crossed, at the calibrated build scale, those whose first entry was through acquisition most often end as leaders or survivors and those that built least often, until builds can start at a larger scale.

The plane's axes have observable counterparts, but locating a firm requires both costs in the model's units. The in-house development budget for a capability corresponds to the build cost, and the licensing or joint-venture fee for the same capability corresponds to the partner cost; both must be scaled to the model's units before a point on the plane has meaning. Contractibility corresponds to how far the capability can be written into a contract or protected by a patent or standard rather than living in tacit know-how, and pace to how fast demand is moving to the new category. These observed quantities support comparisons across firms and settings; where the boundaries fall in a particular industry is an empirical question the model does not settle.

6 Limitations and future research

We separate caveats that qualify specific results from scope conditions of the design. Five caveats travel with the results. First, the concentration metric is tied to the mechanism that generates it, because the Herfindahl index enters the dominant-design term that drives sharpening. The knockout and the decoupled check of Section 4.8 show that concentration comes from the sharpening, but the decoupled level exceeds the F2 band, so how we measure concentration is not fully separated from how it is built. Second, the size of the shakeout is set by calibration: we tuned the exit floor to land inside the stylized-fact bands instead of deriving it from first principles, so the mechanism is causal and switchable, while its size is set jointly by the mechanism and the floor. Third, the erosion of the first-mover advantage with pace appears only with the imitation and obsolescence channels, which are off at baseline; at baseline the gap rises with pace instead (F5). The randomization tests show that most of the gap in the timing experiment survives removing attention-driven selection and strength-based entry order; they do not make erosion a property of the baseline model. Fourth, the homogenization knockout shows that outcome inequality persists, and grows, after initial resources are homogenized; it does not identify the sorting channel, which rests on the tercile comparison of the heterogeneous baseline. Fifth, the multi-region structure

depends on choice noise: as the decision temperature approaches the best-response limit, the regions narrow toward one mode even though firms remain heterogeneous, so the contingency rests on bounded rationality as much as on heterogeneity. Relatedly, build's region on the phase map reflects how often builds are chosen more than how well builders do: scored per entrant, building leads in almost no cells.

The K5 knockout concerns permitted timing: the matched cohorts pass the same attention gate, so the strength-matched gap is the effect of being allowed to enter earlier, not of entering unconditionally (Section 4.3). The early wave also has the target pool to itself until year 12, so part of that gap may reflect first access to targets.

The remaining limits are scope conditions built into the design. Demand is split by capability-weighted shares, with no explicit price and no consumer segments, so differences in willingness to pay and segment-by-segment adoption sit outside the present model (Adner, 2002). Route choice is evaluated one year at a time, not optimized over multi-step sequences, so the option-like logic of staged moves such as ally-then-acquire is not itself a choice the firm makes (Kogut, 1991). The incumbent's dilemma of eating into its own legacy core enters only through the legacy decline that triggers the crossing, not as an explicit loss that entering imposes on the parent (Christensen, 1997). Acquisition is limited by target availability and integration friction; merger review enters only in a robustness run, as a price friction that rises with dominant-design strength, and that run shows the omission from the baseline shapes the maps but not the conclusion that no mode wins everywhere. Alliances add capacity to the market, because an alliance unit's capacity draws on a partner that stays active with its own. Finally, the two calibrated primitives are tuned to one set of stylized-fact bands, not to several industries, so the exact positions of the boundaries on the phase map should be read as illustrative until the parameters are recalibrated against further settings such as semiconductors, enterprise software, or mobility.

7 Conclusion

Capability sourcing under a moving frontier is not a firm-level optimization with a portable answer. Each acquisition uses up a target the next mover could have bought, early commitments raise the barrier later entrants face, and every commitment changes the opportunities the next firm reads. In the model, these interactions shape sourcing comparisons in three ways.

The first concerns opportunities. Rivals' acquisitions change which routes are feasible, and the route mix follows feasibility rather than preference. The second concerns timing. Early entry carries a time-in-market advantage that survives randomized timing and entry order and stays positive at every pace tested; with the imitation and obsolescence channels on, it narrows at the fastest paces only when firms time their own entry. The third concerns the outcome measure. In the plane spanned by the costs of building and partnering, build, partner, and acquire each produce the most leaders in a region, yet per entrant, with acquisition at its baseline cost, acquisition leads in most of the plane, partnering keeps the region where it is cheap, and build's region reflects how often firms build. At the calibrated build scale, acquisition also has the highest leader-or-survivor rate among entering parents, a ranking in which partnering overtakes acquisition when builds start larger. Beyond these three, resources sort firms into routes, while the size of outcome inequality builds up through competition. The feasibility, timing, and inequality findings are each paired with a manipulation of their mechanism (Section 4.8); the ranking among entrants is tested by sweeping the build scale (Section 4.7).

The model does not forecast any one industry. What it helps a reader distinguish is which of three questions a sourcing comparison answers: which routes are feasible given rivals' earlier moves, which route produces the most leaders under given costs, and what end-of-run outcomes parents entering by each route reach. The answers differ, and a recommendation built on one of them should not be read as an answer to the others.

Statements and Declarations

Competing interests. The authors declare no competing financial interests or personal relationships that could have influenced the work reported in this paper.

Funding. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Ethics approval. Not applicable; the study uses no human participants or personal data.

Author contributions. Yingzheng Liu: Conceptualization, Methodology, Software, Formal analysis, Visualization, Writing – original draft. Shun Cao: Methodology, Validation, Formal analysis, Writing – review & editing. Zhen Liu: Methodology, Formal analysis, Validation, Supervision, Project administration, Writing – original draft, Writing – review & editing. Suning Zhu: Conceptualization, Supervision, Writing – review & editing.

Data availability. The agent-based model and all analysis code are implemented as a documented, tested Python package, and every reported result is reproducible from seeded runs whose seed, parameter hash, and code version are recorded with each output. The model code, the analysis scripts, and the result tables behind every reported number are available from the corresponding author on request during review and will be deposited in a public repository on acceptance.

Declaration of generative AI use. During the preparation of this work the authors used a large language model (Anthropic Claude) to assist with copy-editing of the prose and with writing Python plotting code for the figures. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

Adner, R. (2002). When are technologies disruptive? A demand-based view of the emergence of competition. *Strategic Management Journal*, 23(8):667–688.

Adobe Inc. (2019). Form 10-K for the Fiscal Year Ended November 30, 2018. U.S. Securities and Exchange Commission filing. `https://www.sec.gov/Archives/edgar/data/796343/000079634319000019/adbe10kfy18.htm`.

Adobe Inc. (2023). Adobe and Figma Mutually Agree to Terminate Merger Agreement. Press release. `https://news.adobe.com/news/news-details/2023/adobe-and-figma-mutually-agree-to-terminate-merger-agreement`.

Ahuja, G. and Katila, R. (2001). Technological acquisitions and the innovation performance of acquiring firms: a longitudinal study. *Strategic Management Journal*, 22(3):197–220.

Alamalhodaei, A. (2021). Ford and SK Innovation Announce Battery Manufacturing Joint Venture

BlueOvalSK. TechCrunch. https://techcrunch.com/2021/05/20/ford-and-sk-innovation-announce-battery-manufacturing-joint-venture-blueovalsk/.

Anderson, P. and Tushman, M. L. (1990). Technological discontinuities and dominant designs: a cyclical model of technological change. *Administrative Science Quarterly*, 35(4):604–633.

Arthur, W. B. (1989). Competing technologies, increasing returns, and lock-in by historical events. *The Economic Journal*, 99(394):116–131.

Axtell, R. L. and Farmer, J. D. (2025). Agent-based modeling in economics and finance: past, present, and future. *Journal of Economic Literature*, 63(1):197–287.

Bankes, S. (1993). Exploratory modeling for policy analysis. *Operations Research*, 41(3):435–449.

Benlian, A. and Hess, T. (2011). Opportunities and risks of software-as-a-service: findings from a survey of IT executives. *Decision Support Systems*, 52(1):232–246.

Bhargava, S. C. (1989). Generalized Lotka–Volterra equations and the mechanism of technological substitution. *Technological Forecasting and Social Change*, 35(4):319–326.

Bitencourt, C. C., de Oliveira Santini, F., Ladeira, W. J., Santos, A. C., and Teixeira, E. K. (2020). The extended dynamic capabilities model: a meta-analysis. *European Management Journal*, 38(1):108–120.

Blanco-Fernández, D., Leitner, S., and Rausch, A. (2025). Interactions between dynamic team composition and coordination: an agent-based modeling approach. *Review of Managerial Science*, 19(1):1–37.

Bonabeau, E. (2002). Agent-based modeling: methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, 99(suppl 3):7280–7287.

Bourgeois, L. J. (1981). On the measurement of organizational slack. *Academy of Management Review*, 6(1):29–39.

Brenner, T. and Werker, C. (2007). A taxonomy of inference in simulation models. *Computational Economics*, 30(3):227–244.

Brenner, T. and Werker, C. (2009). Policy advice derived from simulation models. *Journal of Artificial Societies and Social Simulation*, 12(4):2.

Brewis, C., Dibb, S., and Meadows, M. (2023). Leveraging big data for strategic marketing: a dynamic capabilities model for incumbent firms. *Technological Forecasting and Social Change*, 190:122402.

Burgelman, R. A. (1983). A process model of internal corporate venturing in the diversified major firm. *Administrative Science Quarterly*, 28(2):223–244.

Capron, L. (1999). The long-term performance of horizontal acquisitions. *Strategic Management Journal*, 20(11):987–1018.

Capron, L. and Mitchell, W. (2009). Selection capability: how capability gaps and internal social frictions affect internal and external strategic renewal. *Organization Science*, 20(2):294–312.

Capron, L. and Mitchell, W. (2012). *Build, Borrow, or Buy: Solving the Growth Dilemma*. Harvard Business Review Press, Boston, MA.

Cassiman, B. and Veugelers, R. (2006). In search of complementarity in innovation strategy: internal R&D and external knowledge acquisition. *Management Science*, 52(1):68–82.

Cefis, E., Marsili, O., and Schenk, H. (2009). The effects of mergers and acquisitions on the firm size distribution. *Journal of Evolutionary Economics*, 19(1):1–20.

Chanda, S. S. and Miller, K. D. (2019). Replicating agent-based models: revisiting March's exploration–exploitation study. *Strategic Organization*, 17(4):425–449.

Christensen, C. M. (1997). *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Harvard Business School Press, Boston, MA.

Christensen, C. M. and Bower, J. L. (1996). Customer power, strategic investment, and the failure of leading firms. *Strategic Management Journal*, 17(3):197–218.

Cohen, W. M. and Levinthal, D. A. (1990). Absorptive capacity: a new perspective on learning and innovation. *Administrative Science Quarterly*, 35(1):128–152.

Cusumano, M. A. (2010). Cloud computing and SaaS as new computing platforms. *Communications of the ACM*, 53(4):27–29.

Cyert, R. M. and March, J. G. (1963). *A Behavioral Theory of the Firm*. Prentice-Hall, Englewood Cliffs, NJ.

Davis, J. P., Eisenhardt, K. M., and Bingham, C. B. (2007). Developing theory through simulation methods. *Academy of Management Review*, 32(2):480–499.

Dawid, H. (2006). Agent-based models of innovation and technological change. In Tesfatsion, L. and Judd, K. L., editors, *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*, chapter 25, pages 1235–1272. Elsevier, Amsterdam.

Dawid, H. and Harting, P. (2012). Capturing firm behavior in agent-based models of industry evolution and macroeconomic dynamics. In Buenstorf, G., editor, *Evolution, Organization and Economic Behavior*, pages 103–130. Edward Elgar, Cheltenham.

Dawid, H., Harting, P., van der Hoog, S., and Neugart, M. (2019). Macroeconomics with heterogeneous agent models: fostering transparency, reproducibility and replication. *Journal of Evolutionary Economics*, 29(1):467–538.

Dawid, H. and Hellmann, T. (2014). The evolution of R&D networks. *Journal of Economic Behavior & Organization*, 105:158–172.

Dawid, H. and Hepp, J. (2022). Distributional effects of technological regime changes: hysteresis, concentration and inequality dynamics. *Review of Evolutionary Political Economy*, 3(1):137–167.

Dierickx, I. and Cool, K. (1989). Asset stock accumulation and sustainability of competitive advantage. *Management Science*, 35(12):1504–1511.

Dixit, A. K. and Pindyck, R. S. (1994). *Investment under Uncertainty*. Princeton University Press, Princeton, NJ.

Dosi, G., Fagiolo, G., and Roventini, A. (2010). Schumpeter meeting Keynes: a policy-friendly model of endogenous growth and business cycles. *Journal of Economic Dynamics and Control*, 34(9):1748–1767.

Dyer, J. H. and Singh, H. (1998). The relational view: cooperative strategy and sources of interorganizational competitive advantage. *Academy of Management Review*, 23(4):660–679.

Eastman Kodak Company (2012). Eastman Kodak Company and Its U.S. Subsidiaries Commence Voluntary Chapter 11 Business Reorganization. Press release, Exhibit 99.1 to Form 8-K. https://www.sec.gov/Archives/edgar/data/0000031235/000119312512016443/d285282dex991.htm.

Eggers, J. P. and Park, K. F. (2018). Incumbent adaptation to technological change: the past, present, and future of research on heterogeneous incumbent response. *Academy of Management Annals*, 12(1):357–389.

Eisenhardt, K. M. and Martin, J. A. (2000). Dynamic capabilities: what are they? *Strategic Management Journal*, 21(10–11):1105–1121.

Epstein, J. M. (1999). Agent-based computational models and generative social science. *Complexity*, 4(5):41–60.

Estrin, J. (2015). Kodak's First Digital Moment. The New York Times, Lens blog. https://archive.nytimes.com/lens.blogs.nytimes.com/2015/08/12/kodaks-first-digital-moment/.

Fagiolo, G., Moneta, A., and Windrum, P. (2007). A critical guide to empirical validation of agent-based models in economics: methodologies, procedures, and open problems. *Computational Economics*, 30(3):195–226.

Fisher, J. C. and Pry, R. H. (1971). A simple substitution model of technological change. *Technological Forecasting and Social Change*, 3:75–88.

Folta, T. B. (1998). Governance and uncertainty: the trade-off between administrative control and commitment. *Strategic Management Journal*, 19(11):1007–1028.

Fontana, R. and Zirulia, L. (2015). "...then came Cisco, and the rest is history": a 'history friendly' model of the Local Area Networking industry. *Journal of Evolutionary Economics*, 25(5):875–899.

Ford Motor Company (2017). Ford Invests in Argo AI, a New Artificial Intelligence Company, in Drive for Autonomous Vehicle Leadership. Press release, Exhibit 99 to Form 8-K. https://www.sec.gov/Archives/edgar/data/0000037996/000003799617000016/exhibit99to8-kdated2x10x20.htm.

FUJIFILM Holdings Corporation (2008). Annual Report 2008 (Year Ended March 31, 2008). Annual report. https://www.annualreports.com/HostedData/AnnualReportArchive/f/fujifilm-holdings-corporation_2008.pdf.

FUJIFILM Holdings Corporation (2011). Annual Report 2011 (Year Ended March 31, 2011). Annual report. https://www.annualreports.com/HostedData/AnnualReportArchive/f/fujifilm-holdings-corporation_2011.pdf.

FUJIFILM Holdings Corporation (2024). 90th Anniversary: History. Corporate website. https://holdings.fujifilm.com/en/about/history/90th.

Gao, D. and Akbaritabar, A. (2022). Using agent-based modeling in routine dynamics research: a quantitative and content analysis of literature. *Review of Managerial Science*, 16(2):521–550.

General Motors (2016). GM to Acquire Cruise Automation to Accelerate Autonomous Vehicle Development. Press release. https://investor.gm.com/news-releases/news-release-details/gm-acquire-cruise-automation-accelerate-autonomous-vehicle.

General Motors (2024). GM to Refocus Autonomous Driving Development on Personal Vehicles. Press release. https://news.gm.com/home.detail.html/Pages/news/us/en/2024/dec/1210-gm.html.

Gilbert, C. G. (2005). Unbundling the structure of inertia: resource versus routine rigidity. *Academy of Management Journal*, 48(5):741–763.

Gilbert, N., Pyka, A., and Ahrweiler, P. (2001). Innovation networks: a simulation approach. *Journal of Artificial Societies and Social Simulation*, 4(3):8.

Gräbner, C. and Hornykewycz, A. (2022). Capability accumulation and product innovation: an agent-based perspective. *Journal of Evolutionary Economics*, 32(1):87–121.

Greve, H. R. (2003). *Organizational Learning from Performance Feedback: A Behavioral Perspective on Innovation and Change*. Cambridge University Press, Cambridge.

Grimm, V., Berger, U., DeAngelis, D. L., Polhill, J. G., Giske, J., and Railsback, S. F. (2010). The ODD protocol: a review and first update. *Ecological Modelling*, 221(23):2760–2768.

Gulati, R. (1995). Does familiarity breed trust? The implications of repeated ties for contractual choice in alliances. *Academy of Management Journal*, 38(1):85–112.

Hanelt, A., Bohnsack, R., Marz, D., and Antunes Marante, C. (2021). A systematic review of the literature on digital transformation: insights and implications for strategy and organizational change. *Journal of Management Studies*, 58(5):1159–1197.

Hannan, M. T. and Freeman, J. (1984). Structural inertia and organizational change. *American Sociological Review*, 49(2):149–164.

Harrison, J. R., Lin, Z., Carroll, G. R., and Carley, K. M. (2007). Simulation modeling in organizational and management research. *Academy of Management Review*, 32(4):1229–1245.

Haspeslagh, P. C. and Jemison, D. B. (1991). *Managing Acquisitions: Creating Value Through Corporate Renewal*. Free Press, New York.

Helfat, C. E., Finkelstein, S., Mitchell, W., Peteraf, M. A., Singh, H., Teece, D. J., and Winter, S. G. (2007). *Dynamic Capabilities: Understanding Strategic Change in Organizations*. Blackwell, Malden, MA.

Helfat, C. E. and Raubitschek, R. S. (2018). Dynamic and integrative capabilities for profiting from innovation in digital platform-based ecosystems. *Research Policy*, 47(8):1391–1399.

Henderson, R. M. and Clark, K. B. (1990). Architectural innovation: the reconfiguration of existing product technologies and the failure of established firms. *Administrative Science Quarterly*, 35(1):9–30.

Jacobides, M. G., Cennamo, C., and Gawer, A. (2018). Towards a theory of ecosystems. *Strategic Management Journal*, 39(8):2255–2276.

Jovanovic, B. (1982). Selection and the evolution of industry. *Econometrica*, 50(3):649–670.

Kale, P. and Singh, H. (2007). Building firm capabilities through learning: the role of the alliance learning process in alliance capability and firm-level alliance success. *Strategic Management Journal*, 28(10):981–1000.

Karim, S. (2006). Modularity in organizational structure: the reconfiguration of internally developed and acquired business units. *Strategic Management Journal*, 27(9):799–823.

Katz, M. L. and Shapiro, C. (1985). Network externalities, competition, and compatibility. *American Economic Review*, 75(3):424–440.

King, D. R., Dalton, D. R., Daily, C. M., and Covin, J. G. (2004). Meta-analyses of post-acquisition performance: indications of unidentified moderators. *Strategic Management Journal*, 25(2):187–200.

King, D. R., Wang, G., Samimi, M., and Cortes, A. F. (2021). A meta-analytic integration of acquisition performance prediction. *Journal of Management Studies*, 58(5):1198–1236.

Klepper, S. (1996). Entry, exit, growth, and innovation over the product life cycle. *American Economic Review*, 86(3):562–583.

Klepper, S. (1997). Industry life cycles. *Industrial and Corporate Change*, 6(1):145–182.

Klepper, S. and Simons, K. L. (1997). Technological extinctions of industrial firms: an inquiry into their nature and causes. *Industrial and Corporate Change*, 6(2):379–460.

Klepper, S. and Simons, K. L. (2000). The making of an oligopoly: firm survival and technological change in the evolution of the U.S. tire industry. *Journal of Political Economy*, 108(4):728–760.

Kogut, B. (1991). Joint ventures and the option to expand and acquire. *Management Science*, 37(1):19–33.

Korosec, K. (2022). Ford, VW-backed Argo AI Is Shutting Down. TechCrunch. `https://techcrunch.com/2022/10/26/ford-vw-backed-argo-ai-is-shutting-down/`.

Kwakkel, J. H. and Pruyt, E. (2013). Exploratory modeling and analysis, an approach for model-based foresight under deep uncertainty. *Technological Forecasting and Social Change*, 80(3):419–431.

Lane, P. J. and Lubatkin, M. (1998). Relative absorptive capacity and interorganizational learning. *Strategic Management Journal*, 19(5):461–477.

Lardinois, F. (2013). Adobe Goes All In With Subscription-Based Creative Cloud, Will Stop Selling Regular Licenses. TechCrunch. `https://techcrunch.com/2013/05/06/adobe-goes-all-in-with-subscription-based-creative-cloud-will-stop-selling-regular-cs-licenses-shrink-wrapped-boxes/`.

Lieberman, M. B. and Montgomery, D. B. (1988). First-mover advantages. *Strategic Management Journal*, 9(S1):41–58.

Lieberman, M. B. and Montgomery, D. B. (1998). First-mover (dis)advantages: retrospective and link with the resource-based view. *Strategic Management Journal*, 19(12):1111–1125.

Malerba, F., Nelson, R., Orsenigo, L., and Winter, S. (1999). 'History-friendly' models of industry evolution: the computer industry. *Industrial and Corporate Change*, 8(1):3–40.

Malerba, F., Nelson, R. R., Orsenigo, L., and Winter, S. G. (2016). *Innovation and the Evolution of Industries: History-Friendly Models*. Cambridge University Press, Cambridge.

Mallapragada, G., Bommaraju, R., Kumar, A., and Pedada, K. (2025). To acquire or to ally? The impact of strategic emphasis on governance mode choice. *Journal of Marketing Research*, 62(4):645–663.

Marasco, A., Picucci, A., and Romano, A. (2016). Market share dynamics using Lotka–Volterra models. *Technological Forecasting and Social Change*, 105:49–62.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In Zarembka, P., editor, *Frontiers in Econometrics*, pages 105–142. Academic Press, New York.

Metcalfe, J. S. (1994). Competition, Fisher's principle and increasing returns in the selection process. *Journal of Evolutionary Economics*, 4(4):327–346.

Microsoft (2010). Windows Azure General Availability. The Official Microsoft Blog. `https://blogs.microsoft.com/blog/2010/02/01/windows-azure-general-availability/`.

Microsoft (2016). Microsoft to Acquire LinkedIn. Press release. `https://news.microsoft.com/source/2016/06/13/microsoft-to-acquire-linkedin/`.

Microsoft (2018). Microsoft to Acquire GitHub for $7.5 Billion. Press release. `https://news.microsoft.com/source/2018/06/04/microsoft-to-acquire-github-for-7-5-billion/`.

Miller, K. D. and Lin, S.-J. (2022). Strategic issue diagnosis by top management teams: a multiple-agent model. *Strategic Organization*, 20(3):600–626.

Morris, S. A. and Pratt, D. (2003). Analysis of the Lotka–Volterra competition equations as a technological substitution model. *Technological Forecasting and Social Change*, 70(2):103–133.

Müller, B., Bohn, F., Dreßler, G., Groeneveld, J., Klassert, C., Martin, R., Schlüter, M., Schulze, J., Weise, H., and Schwarz, N. (2013). Describing human decisions in agent-based models – ODD+D, an extension of the ODD protocol. *Environmental Modelling & Software*, 48:37–48.

Nelson, R. R. and Winter, S. G. (1982). *An Evolutionary Theory of Economic Change*. Belknap Press of Harvard University Press, Cambridge, MA.

Nohria, N. and Gulati, R. (1996). Is slack good or bad for innovation? *Academy of Management Journal*, 39(5):1245–1264.

Oracle (2016). Oracle Buys NetSuite. Press release. `https://www.oracle.com/corporate/pressrelease/oracle-buys-netsuite-072816.html`.

Penrose, E. T. (1959). *The Theory of the Growth of the Firm*. Oxford University Press, Oxford.

Perkins, G. and Murmann, J. P. (2018). What does the success of Tesla mean for the future dynamics in the global automobile sector? *Management and Organization Review*, 14(3):471–480.

Pisano, G. P. (1990). The R&D boundaries of the firm: an empirical analysis. *Administrative Science Quarterly*, 35(1):153–176.

Pitelis, C. N., Teece, D. J., and Yang, H. (2024). Dynamic capabilities and MNE global strategy: a systematic literature review-based novel conceptual framework. *Journal of Management Studies*, 61(7):3295–3326.

Puranam, P., Singh, H., and Zollo, M. (2006). Organizing for innovation: managing the coordination-autonomy dilemma in technology acquisitions. *Academy of Management Journal*, 49(2):263–280.

Railsback, S. F. and Grimm, V. (2019). *Agent-Based and Individual-Based Modeling: A Practical Introduction*. Princeton University Press, Princeton, NJ, second edition.

Rand, W. and Rust, R. T. (2011). Agent-based modeling in marketing: guidelines for rigor. *International Journal of Research in Marketing*, 28(3):181–193.

Robertson, D. A. (2019). *Agent-Based Strategizing*. Elements in Business Strategy. Cambridge University Press, Cambridge.

Rohrbeck, R. and Kum, M. E. (2018). Corporate foresight and its impact on firm performance: a longitudinal analysis. *Technological Forecasting and Social Change*, 129:105–116.

Salesforce (2018). Salesforce Signs Definitive Agreement to Acquire MuleSoft. Press release. https://www.salesforce.com/news/press-releases/2018/03/20/salesforce-signs-definitive-agreement-to-acquire-mulesoft/.

Salesforce (2019). Salesforce Signs Definitive Agreement to Acquire Tableau. Press release. https://www.salesforce.com/news/press-releases/2019/06/10/salesforce-signs-definitive-agreement-to-acquire-tableau/.

Salesforce (2021). Salesforce Completes Acquisition of Slack. Press release, Exhibit 99.1 to Form 8-K. https://www.sec.gov/Archives/edgar/data/1108524/000119312521220235/d159530dex991.htm.

Salesforce and Slack Technologies (2020). Salesforce Signs Definitive Agreement to Acquire Slack. Press release, Exhibit 99.2 to Slack Technologies Form 8-K. https://www.sec.gov/Archives/edgar/data/1764925/000119312520307385/d59995dex992.htm.

Saltelli, A., Annoni, P., Azzini, I., Campolongo, F., Ratto, M., and Tarantola, S. (2010). Variance based sensitivity analysis of model output: design and estimator for the total sensitivity index. *Computer Physics Communications*, 181(2):259–270.

SAP SE (2014). SAP to acquire Concur, expanding the world's largest business network. Press release, Exhibit 99.1 to Form 6-K. https://www.sec.gov/Archives/edgar/data/0001000184/000119312514345861/d791851dex991.htm.

SAP SE (2018). SAP SE to acquire Qualtrics International Inc. Press release, Exhibit 99.1 to Form 6-K. https://www.sec.gov/Archives/edgar/data/0001000184/000110465918067900/a18-40032_2ex99d1.htm.

SAP SE (2019). SAP completes acquisition of Qualtrics. Press release (PR Newswire). https://www.prnewswire.com/news-releases/sap-completes-acquisition-of-qualtrics-300782755.html.

Schilke, O. and Helfat, C. E. (2025). Unlocking dynamic capabilities: pathways for empirical research. *Journal of Management Scientific Reports*, 3(2):71–87.

Schilke, O., Hu, S., and Helfat, C. E. (2018). Quo vadis, dynamic capabilities? A content-analytic review of the current state of knowledge and recommendations for future research. *Academy of Management Annals*, 12(1):390–439.

Schwarz, J. O., Rohrbeck, R., and Wach, B. (2020). Corporate foresight as a microfoundation of dynamic capabilities. *Futures & Foresight Science*, 2(2):e28.

Silverberg, G., Dosi, G., and Orsenigo, L. (1988). Innovation, diversity and diffusion: a self-organisation model. *The Economic Journal*, 98(393):1032–1054.

Stinchcombe, A. L. (1965). Social structure and organizations. In March, J. G., editor, *Handbook of Organizations*, pages 142–193. Rand McNally, Chicago.

Strohmaier, R. and Rainer, A. (2016). Studying general purpose technologies in a multi-sector framework: the case of ICT in Denmark. *Structural Change and Economic Dynamics*, 36:34–49.

Suarez, F. F., Grodal, S., and Gotsopoulos, A. (2015). Perfect timing? Dominant category, dominant design, and the window of opportunity for firm entry. *Strategic Management Journal*, 36(3):437–448.

Suarez, F. F. and Lanzolla, G. (2007). The role of environmental dynamics in building a first mover advantage theory. *Academy of Management Review*, 32(2):377–392.

Suarez, F. F. and Utterback, J. M. (1995). Dominant designs and the survival of firms. *Strategic Management Journal*, 16(6):415–430.

Teece, D. J. (2007). Explicating dynamic capabilities: the nature and microfoundations of (sustainable) enterprise performance. *Strategic Management Journal*, 28(13):1319–1350.

Teece, D. J., Pisano, G., and Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7):509–533.

ten Broeke, G., van Voorn, G., and Ligtenberg, A. (2016). Which sensitivity analysis method should I use for my agent-based model? *Journal of Artificial Societies and Social Simulation*, 19(1):5.

Tripsas, M. and Gavetti, G. (2000). Capabilities, cognition, and inertia: evidence from digital imaging. *Strategic Management Journal*, 21(10–11):1147–1161.

Tushman, M. L. and Anderson, P. (1986). Technological discontinuities and organizational environments. *Administrative Science Quarterly*, 31(3):439–465.

Verhoef, P. C., Broekhuizen, T., Bart, Y., Bhattacharya, A., Dong, J. Q., Fabian, N., and Haenlein, M. (2021). Digital transformation: a multidisciplinary reflection and research agenda. *Journal of Business Research*, 122:889–901.

Veugelers, R. (1997). Internal R&D expenditures and external technology sourcing. *Research Policy*, 26(3):303–315.

Vial, G. (2019). Understanding digital transformation: a review and a research agenda. *The Journal of Strategic Information Systems*, 28(2):118–144.

Villalonga, B. and McGahan, A. M. (2005). The choice among acquisitions, alliances, and divestitures. *Strategic Management Journal*, 26(13):1183–1208.

Volkswagen Group (2020). Volkswagen Increases Stake in QuantumScape. Press release no. 124/2020. `https://www.volkswagen-group.com/en/press-releases/volkswagen-increases-stake-in-quantumscape-17068`.

Warner, K. S. R. and Wäger, M. (2019). Building dynamic capabilities for digital transformation: an ongoing process of strategic renewal. *Long Range Planning*, 52(3):326–349.

Wayland, M. (2019). GM, LG Chem to form $2.3 billion joint venture for battery cell production. CNBC. `https://www.cnbc.com/2019/12/05/gm-lg-to-form-2point3-billion-joint-venture-for-battery-cell-production.html`.

Werker, C. (2003). Innovation, market performance, and competition: lessons from a product life cycle model. *Technovation*, 23(4):281–290.

Werker, C. and Brenner, T. (2004). Empirical calibration of simulation models. Papers on Economics and Evolution 0410, Max Planck Institute for Research into Economic Systems, Jena.

Werker, C., Korzinov, V., and Cunningham, S. (2019). Formation and output of collaborations: the role of proximity in German nanotechnology. *Journal of Evolutionary Economics*, 29(2):697–719.

West, J. and Bogers, M. (2014). Leveraging external sources of innovation: a review of research on open innovation. *Journal of Product Innovation Management*, 31(4):814–831.

Williamson, O. E. (1985). *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*. Free Press, New York.

Windrum, P., Fagiolo, G., and Moneta, A. (2007). Empirical validation of agent-based models: alternatives and prospects. *Journal of Artificial Societies and Social Simulation*, 10(2):8.

Winter, S. G., Kaniovski, Y. M., and Dosi, G. (2003). A baseline model of industry evolution. *Journal of Evolutionary Economics*, 13(4):355–383.

Wu, J., Ohya, K., and Sekiguchi, T. (2024). Applications of agent-based modeling and simulation in organization management: a quarter-century review through bibliometric mapping (1998–2022). *Computational and Mathematical Organization Theory*, 30(1):1–31.

Zachary, M. A., Gianiodis, P. T., Payne, G. T., and Markman, G. D. (2015). Entry timing: enduring lessons and future directions. *Journal of Management*, 41(5):1388–1415.

Zollo, M. and Winter, S. G. (2002). Deliberate learning and the evolution of dynamic capabilities. *Organization Science*, 13(3):339–351.

Appendix A: ODD+D protocol

This appendix documents the model in full, following the ODD protocol and its +Decision extension for human decision-making in agent-based models (Grimm et al., 2010; Müller et al., 2013). Symbols follow the main text (Table 3); every parameter, with its governed range and literature provenance, is listed in Appendix B. The model is implemented as a tested Python package; all stochasticity in a run flows from one seeded generator, and every output carries its seed, parameter hash, and code commit, so each reported number is reproducible.

A.1 Purpose and patterns

The model's purpose is forward reasoning, not data fitting, about how an incumbent population sources capability (build, partner, acquire, or wait) across a legacy-to-emerging market transition, and about the market architecture those choices produce. It is used for *generative foresight*: directional structure derived from evidence-grounded rules rather than from a historical panel. Credibility rests on two supports. The first is pattern-oriented validation against six documented stylized facts (F1 diffusion-plus-shakeout; F2 rising concentration; F3 alliance-under-uncertainty; F4 acquisitions for distant capability gaps; F5 pace-contingent first-mover advantage; F6 alliance-unit fragility), of which two are calibration targets (F1, F2), one untargeted check passes (F4), two are not reproduced (F3, F5), and one is unscored (F6) (Table 10). The second is knockout identification, in which each generative, confound, or decomposition headline must collapse or fall sharply, persist, or reveal its channel when its candidate mechanism is removed, randomized, or decomposed; the entrant ranking is examined instead by a build-scale sweep.

A.2 Entities, state variables, and scales

Three entity classes inhabit two coupled markets, a declining legacy market L and a growing emerging market E. The first two are firms in their own right: *legacy parents* are the decision-makers, and *emerging independents* enter endogenously and serve simultaneously as entrants, acquisition targets, and alliance partners. The third class is the *affiliate*, the unit an incumbent creates to enter, which takes the form of an internal build seeded from parent capability, an acquired unit (a former independent), or an alliance unit held at equity stake θ. Every firm or affiliate carries the capability stocks $K, T, A, L \in [0, 1]$ (knowledge, technology, adaptability, legitimacy), cash F, capacity C, captured demand q, share s, an age, gestation timers (ramp, integration, coordination-drag, newness), and profit/finance/share/technology histories (Table 3). Market-level state is the year t, total demand D_t, emerging share m_t, the legacy and emerging frontiers T_t^L, T_t^E, the signed frontier gap, emerging concentration HHI_t^E and mean quality $\bar{Q}_t^E$, dominant-design strength S_t, entry barrier B_t, lifecycle stage, stage uncertainty, the innovation shock, and the disruption flag. Time is discrete and annual, $t = 0, 1, \ldots, T$ with horizon $T = 60$ (fifty in Experiment 3; the robustness designs state their own horizons in Table 8).

A.3 Process overview and scheduling

Each year executes the same ordered schedule. (1) The market update advances: it draws an innovation shock, reads the frontiers and concentration off the previous year's population, computes dominant-design strength S_t, moves the emerging share m_t by the replicator rule, and then updates

the entry barrier B_t, the life-cycle stage and uncertainty, the disruption flag, and the demand split. (2) Legacy incumbents, in randomized order, run the sense → screen → seize pipeline; a firm that decides to act creates the corresponding affiliate and its target/partner is consumed or engaged. (3) Emerging independents enter de novo by a stage-modulated Poisson process whose intensity opens with the category and tapers as the dominant design consolidates, the window-of-opportunity pattern of industry evolution (Suarez et al., 2015; Klepper, 1996). (4) Demand is allocated, legacy market first in proportion to position, then emerging with dominant-design-gated increasing-returns sharpening. (5) Operations settle: profit, external finance, capacity adjustment, and capability learning, per active firm, followed by inter-firm transfers (spillovers, dividends, parent support). (6) Ages and timers advance, and entity-specific selection removes firms that cannot sustain themselves. Per-year metrics and the decision panel are then recorded.

A.4 Design concepts

Basic principles. The behavioral architecture is dynamic capability theory (sensing, seizing, transforming) joined to the resource-pathways logic of Capron and Mitchell; market dynamics use a replicator (competitive-substitution) law and a single bounded increasing-returns primitive. **Emergence.** Concentration, the shakeout, the mode-dominant regions, the leader-or-survivor ranking, and outcome stratification are outcomes, not parameters; each is held to a knockout, a sweep, or both. **Adaptation and objectives.** A firm scores each feasible route by an expected position value net of cost and chooses by softmax (a boundedly rational random-utility rule), so behavior approximates, but does not assume, payoff maximization. **Learning.** Capability stocks accumulate by investment gated through absorptive capacity, with route-specific transfer, spillover, and obsolescence. **Prediction.** Firms are myopic: route values use a short planning horizon and current state, not forecasts. **Sensing.** Attention to the emerging market is an evidence-weighted function of adverse-feedback and incumbency signals. **Interaction.** Firms interact through competition for finite demand, depletion of the target pool by acquisition, bilateral alliance acceptance, and capability spillovers. **Stochasticity.** Randomness enters through the correlated initialization, annual innovation shocks, de-novo entry counts, and the softmax choice draw; all are seeded. **Observation.** Per-year market and population metrics, a per-decision panel, and end-of-run outcome classifications are logged. **Decision-making.** Appendix C records the theoretical anchor of each decision rule; Section 3.4 gives the alternatives the agent weighs and how the environment and other agents enter the choice.

A.5 Initialization

Firms are drawn continuously within overlapping seed size classes from a single latent maturity factor with signed, evidence-governed loadings; decision rules use the realized resources rather than the class labels. The draw is specified in Eq. 1 and the class marginals and loadings in Table 2 of Section 3.2. Setting $\rho = 0$ would remove the within-class correlation, and setting $\delta = 0$ as well would remove the seed-class mean shifts; the reported runs retain $\rho = 0.6$. The K4 knockout instead sets every firm to the attribute means. A post-hoc resource quantile, the summed z-scores of $\log(1+F)$, K, T, L, and $\log(1+C)$ cut into terciles, labels incumbents for description only and never enters any decision.

A.6 Input data

The model uses no external input data streams. Every governed parameter is a literature-grounded or normalization value drawn from the metadata of Appendix B. The reported analyses vary a subset of them, marked in Table 14; the rest, including the transfer and dividend rates, patience windows, parent support, external finance, and capacity adjustment, are held at their baseline values. Some structural constants are fixed in code, not governed, and appear only in the main-text equations, including the stage entry multipliers of Eq. 9, the absorptive-capacity weights of Eq. 20, the capability weights inside the installed-base quality proxy of Eq. 4, the acceptance steepness of Eq. 13, the scaling constants inside the sense signals of Eq. 10, the half-barrier admission rule, the affiliate cash seed of half the route cost, and the floors on initial age and cash stated with Eq. 1.

A.7 Submodels

Every submodel is specified in full, with its functional form and baseline coefficients, in Section 3 of the main text; nothing is held back for this appendix. Table 12 maps each ODD+D submodel to the equations that define it, so that a reader following the protocol can locate each rule. The per-parameter anchors, governed ranges, and governance classes are in Table 14, and the design–evidence map of Appendix C (Table 15) records the evidence behind each functional form.

Table 12: ODD+D submodels and where each is specified in the main text.

Submodel	Main-text specification
Market update (demand split, competitive substitution, frontiers, lock-in indices, stage and uncertainty, disruption flag)	Eqs. 2–6, Section 3.3
Demand allocation with increasing returns	Eqs. 7–8
De novo entry of emerging independents	Eq. 9
Sense (attention and search pressure)	Eqs. 10–11
Screen (fit and bilateral acceptance)	Eqs. 12–13
Seize (route values, costs, feasibility, softmax choice)	Eqs. 14–17
Affiliate creation and gestation timers	Eq. 18
Transform (learning, imitation, obsolescence)	Eq. 20
Operations, finance, capacity, inter-firm transfers	Eq. 19; transfers in the paragraph "Transfers and parent support" of Sec. 3.5
Select (exit rules by entity and route)	Section 3.5, "Exit (selection)"
Outcome classification and θ-weighted portfolio accounting	Section 3.6, "Outcome class"

Appendix B: parameter inventory

Table 13 lists the parameters that carry the experiments, with their calibrated baseline values, governed ranges, and the reported analyses that vary each. The complete inventory of all parameters, with governance class, governed ranges, and literature provenance, is given in Table 14, auto-generated from the machine-readable metadata that drives the model (`src/params.py`); each one-factor sweep is centered on the baseline recorded there, and the Sobol design draws its ranges from the same metadata.

Table 13: Key parameters: calibrated baseline, governed range, and the reported analyses that vary each.

Role	Parameter	Baseline	Governed range	Varied in
Build cost	`build_fixed_cost`	9.5	4.0–16.0	Exp. 3
Partner cost	`partnership_fixed_cost`	7.0	3.0–12.0	Exp. 3
Acquisition price multiple	`acquisition_price_multiple`	2.5	1.0–4.5	Exp. 3; Sec. 4.8
Pace (substitution rate a)	`base_adoption_rate`	3.4	0.0–4.0	Exps. 1, 2 (to 4.6), 3, 6 (to 4.5)
Target-pool liquidity	`entry_poisson_base`	3.0	1.0–8.0	Exps. 1, 3, 4; K1
Partner reliability	`partnership_accept_threshold`	0.60	0.45–0.80	Exps. 1, 3
Integration friction	`acquisition_integration_loss`	0.20	0.05–0.45	Exps. 1, 3
Contractibility	`opportunity_contractibility`	0.50	0.10–0.90	Exps. 3, 6; knockout
Emergence uncertainty	`uncertainty_emergence`	0.30	0.15–0.45	Exps. 1, 3, 6
Increasing-returns gain	`concentration_gain`	7.0	0.0–8.0	knockout; Sec. 4.8
Sub-scale exit floor	`exit_share_floor`	0.01	0.001–0.01	knockout; Sec. 4.8
Decision temperature	`decision_temperature`	0.33	0.20–0.50	Sec. 4.8 (0.05–0.66)
Resource correlation	`resource_correlation_strength`	0.60	0.00–1.00	not varied

The complete governed parameter inventory follows in Table 14, generated from the parameter metadata in `src/params.py`; for the partner option-value weight, the Range column gives the ±50% range used in the perturbation; the wait option-value weight is not varied. Each parameter carries a governance class, its calibrated baseline, its governed (allowable) range, and its grounding; numerical parameters marked $*$ are varied in at least one reported analysis, and the ranges of the others are allowable ranges, not sweeps; switches (class S) are unmarked, and for them the last column names the intervention each implements. The fifty-percent perturbation of Section 4.8 can reach beyond the listed ranges for the sharpening gain, the exit floor, base obsolescence, and the scale-economy exponent, within validity bounds. The class codes are **N** normalization (free, fixed to a convenient scale), **B** behavioral (evidence-grounded), **E** environmental/structural treatment (the candidates for the phase-map axes), **A** accounting (symmetric across routes or explicitly justified), and **S** switch (a knockout or design toggle, off or symmetric by default).

Table 14: Complete governed parameter inventory, generated from the parameter metadata in `src/params.py`.

Parameter	Base	Cls	Range	Grounding / note
`total_category_demand`	100	N	n/a	base total demand D_0
`category_growth_rate`	0.002	N	n/a	keep aggregate envelope simple
`legacy_exogenous_decline`	0.004	E	0–0.015	annual compounding decay of legacy share, $(1-\delta_L)^t$
`base_adoption_rate`*	3.4	E	0–4 (to 4.6 in Exp. 2, 4.5 in Exp. 6)	Replicator / Lotka-Volterra substitution (Bhargava 1989; Morris & Pratt 2003; Marasco et al. 2016)
`tech_gap_adoption_weight`*	1	E	0–2	Fisher & Pry 1971; Adner 2002 (relative-merit substitution)
`dominant_design_adoption_weight`	0.3	E	0–1	Anderson & Tushman 1990; Arthur 1989 (increasing returns)

continued on next page

Table 14, continued

Parameter	Base	Cls	Range	Grounding / note
intrinsic_emerging_advantage	0	E	−0.1 to 0.1	baseline relative attractiveness (replicator drift)
dominant_design_adoption_floor	0.4	E	0.2–0.6	DD strength must exceed this before it accelerates adoption
initial_emerging_share	0.06	E	0.02–0.2	m_0, initial emerging share
max_emerging_share	0.92	E	0.8–0.99	upper clamp on emerging share; <1 leaves a legacy floor
init_emerging_frontier*	0.62	E	0.3–0.85	design choice (Sec. 3.2)
init_legacy_frontier*	0.45	E	0.3–0.85	design choice (Sec. 3.2)
init_frontier_sd	0.12	E	0.05–0.25	dispersion of initial firm technology around the frontier means
innovation_shock_probability*	0.05	E	0–0.12	annual shock probability
innovation_shock_size	0.06	E	0–0.15	base shock magnitude ($\times U(0.7, 1.3)$)
uncertainty_emergence*	0.3	E	0.15–0.45	stage uncertainty: emergence
uncertainty_growth*	0.22	E	0.1–0.35	stage uncertainty: growth
uncertainty_maturity	0.14	E	0.08–0.25	stage uncertainty: maturity
uncertainty_late	0.12	E	0.06–0.2	stage uncertainty: late
legacy_uncertainty_base	0.09	E	0.04–0.18	baseline legacy-market uncertainty
stage_cut_emergence	0.15	E	0.08–0.25	emergence while m_t is below this
stage_cut_growth	0.45	E	0.3–0.6	growth while m_t is below this
stage_cut_maturity	0.75	E	0.6–0.9	maturity while m_t is below this, late thereafter
dominant_design_share_weight*	0.35	E	0.2–0.5	ω_1 on m_t
dominant_design_frontier_weight*	0.35	E	0.2–0.5	ω_2 on emerging frontier
dominant_design_concentration_weight*	0.3	E	0.15–0.45	ω_3 on emerging HHI
barrier_installed_base_weight	0.35	E	0.2–0.5	η_1 on m_t
barrier_quality_weight	0.35	E	0.2–0.5	η_2 on installed-base quality
barrier_concentration_weight	0.3	E	0.15–0.45	η_3 on emerging HHI
disruption_threshold_adoption	0.3	E	0.2–0.5	m_D: emerging-share gate
disruption_threshold_frontier	0.65	E	0.5–0.85	θ_F: emerging-frontier gate
disruption_frontier_margin	0.08	E	0–0.2	Δ_F: margin by which the emerging frontier must beat the legacy frontier
dominant_design_threshold	0.45	E	0.3–0.65	θ_S: design-strength gate
n_large*	14	E	0–40	number of large incumbents
n_medium*	36	E	0–80	number of medium incumbents
n_small*	72	E	0–140	number of small incumbents
size_class_separation*	0.8	E	0–1.5	continuous size composition (latent-factor mean shift)
emerging_seed_count_low	6	E	3–12	min initial emerging seeds
emerging_seed_count_high	10	E	6–20	max initial emerging seeds
entry_poisson_base*	3	E	1–8	industry-evolution entry rates
startup_entry_persistent_opportunity	0.08	E	0–0.2	persistent floor on de novo entry opportunity
resource_correlation_strength	0.6	B	0–1	Rand & Rust 2011 (sensitivity on initial conditions); Hannan & Freeman 1984
resource_spread_scale	1	B	0.5–1.5	scales all marginal SDs (heterogeneity magnitude); does not affect the K4 mean-collapse
search_threshold	0.52	B	0.45–0.62	Cyert & March 1963; Greve 2003
aspiration_window	3	B	2–5	Greve 2003
attention_gain	2.2	B	1–3.5	g_A: attention sigmoid steepness
attention_threshold	0.25	B	0.1–0.4	θ_A: attention offset
performance_gap_weight	1.1	B	0.55–1.65	Cyert & March 1963
share_decline_weight	0.7	B	0.35–1.05	β_2 share decline
tech_gap_weight	0.95	B	0.48–1.43	β_3 frontier gap
disruption_weight	0.75	B	0.38–1.13	β_4 disruption flag
prior_exposure_attention_bonus	0.2	B	0.1–0.3	β_5 prior-exposure bonus
customer_power_weight	0.7	B	0.35–1.05	Christensen & Bower 1996; Gilbert 2005

continued on next page

Table 14, continued

Parameter	Base	Cls	Range	Grounding / note
`age_inertia_weight`	0.18	B	0.09–0.27	Hannan & Freeman 1984
`attention_inertia_weight`	0.55	B	0.28–0.83	Tripsas & Gavetti 2000
`search_pressure_attention_weight`	0.45	B	0.25–0.65	weight on attention in search pressure (Eq. 11)
`search_pressure_slack_weight`	0.35	B	0.2–0.5	weight on slack in search pressure (Eq. 11)
`search_pressure_adaptability_weight`	0.2	B	0.1–0.3	weight on adaptability in search pressure (Eq. 11)
`slack_peak`	0.35	B	0.2–0.5	Bourgeois 1981; Nohria & Gulati 1996
`slack_width`	0.35	B	0.2–0.5	inverted-U slack width
`cash_normalizer`*	30	B	15–60	κ_F in slack $\varsigma_i = F_i/(F_i + \kappa_F)$
`entry_decision_cooldown_years`	3	B	1–5	min yrs between a parent's entry decisions
`preemption_attention_coef`	0.15	B	0–0.3	sub-threshold preemption proportional to attention
`preemption_disruption_coef`	0.05	B	0–0.15	sub-threshold preemption proportional to disruption
`decision_temperature`*	0.33	B	0.2–0.5	McFadden 1974 (random-utility / softmax)
`planning_horizon`	3	B	1–6	planning horizon h in the gestation discount $\varrho(g)$
`decision_cost_weight`	1	B	0.3–2	λ on the cost ratio in $V_r - \lambda c_r$
`decision_cost_kappa`	1	B	0.5–3	κ in the cost-to-cash ratio $\text{cost}_r/(F_i + \kappa)$
`no_action_utility_bias`	0.3	B	0–0.6	baseline value of the WAIT option
`route_opportunity_threshold`	0.35	B	0.2–0.5	symmetry principle (Sec. 3.4)
`partnership_option_value_weight`*	0.45	B	0.225–0.675	Kogut 1991; Folta 1998 (alliance as real option)
`wait_option_value_weight`	0.05	B	0.025–0.075	Kogut 1991 (option value of waiting)
`wait_slack_value`	0.3	B	0.15–0.45	value of slack retained by WAITing
`irreversibility_penalty`*	0.6	B	0.2–0.8	Folta 1998; Dixit & Pindyck 1994; Kogut 1991 (real options under uncertainty)
`build_fixed_cost`*	9.5	E	4–16	Burgelman 1983; Capron & Mitchell 2012
`build_capacity_share`*	0.22	E	0.1–0.4	Burgelman 1983 (internal venture seeding)
`build_legitimacy_start`	0.28	A	0.15–0.45	build unit's starting legitimacy (low, reflecting the liability of newness)
`partnership_fixed_cost`*	7	E	3–12	Gulati 1995
`acquisition_price_multiple`*	2.5	E	1–4.5	Haspeslagh & Jemison 1991
`acquisition_fixed_cost`	5	E	2–9	acquisition base cost
`acquisition_concentration_penalty`*	0	E	0–3	ξ: merger-review and target-scarcity friction, robustness runs only (Haspeslagh & Jemison 1991)
`acquisition_integration_loss`*	0.2	A	0.05–0.45	Ahuja & Katila 2001; Puranam et al. 2006
`build_ramp_years`	3	E	1–6	build ramp gestation
`partnership_years`	3	E	2–6	alliance coordination-drag window
`integration_years`	3	E	1–6	post-acquisition integration window
`partnership_accept_threshold`*	0.6	E	0.45–0.8	Dyer & Singh 1998; Gulati 1995
`leakage_risk_weight`	0.3	B	0.15–0.45	Dyer & Singh 1998 (appropriation/leakage hazard)
`alliance_equity_stake`*	0.5	A	0.3–1	explicit equity stake (Sec. 3.4)
`alliance_coordination_penalty`*	0.2	A	0.05–0.45	Dyer & Singh 1998; Kale & Singh 2007 (shared-control coordination/leakage cost)
`opportunity_contractibility`*	0.5	E	0.1–0.9	Williamson 1985 (asset specificity); Capron & Mitchell 2012; Mallapragada et al. 2025

continued on next page

Table 14, continued

Parameter	Base	Cls	Range	Grounding / note
contractibility_partner_gain*	1.2	B	0–2	Capron & Mitchell 2012; Mallapragada et al. 2025 (value appropriation favors acquisition; the contractibility reading is ours)
contractibility_internalize_gain*	0.8	B	0–2	Capron & Mitchell 2012 (internalize when non-contractible); Williamson 1985
leader_threshold_quantile*	0.75	A	0.6–0.9	top-quantile leader cutoff, applied identically across routes; varied in the specification variants
alliance_dividend_share_to_legacy	0.3	A	0.15–0.45	share of alliance profit to legacy sponsor
alliance_dividend_share_to_partner	0.2	A	0.1–0.35	share of alliance profit to partner
acquisition_dividend_share_to_parent	0.5	A	0.3–0.7	share of acquired-unit profit to parent
parent_transfer_rate	0.18	A	0.09–0.27	Karim 2006 (venture inherits parent K)
reverse_learning_rate	0.08	A	0.04–0.12	Capron & Mitchell 2012
alliance_spillover_rate	0.1	A	0.05–0.15	Lane & Lubatkin 1998
acquisition_spillover_rate	0.1	A	0.05–0.15	Capron 1999
strategic_survival_demand_floor	0.35	A	0.2–0.6	min portfolio demand to count as survivor
venture_patience_years	6	B	3–9	Stinchcombe 1965
partner_patience_years	5	B	3–8	alliance patience window
acquire_patience_years	5	B	3–8	acquisition patience window
liability_of_newness_years	2	B	1–4	Stinchcombe 1965
newness_penalty	0.12	B	0.06–0.18	demand-weight penalty during newness
base_rd_share	0.08	B	0.04–0.12	R&D propensity (fraction of cash)
base_capex_share	0.05	B	0.025–0.075	capex propensity
knowledge_invest_efficiency	0.11	B	0.055–0.165	knowledge gain per scaled R&D
technology_invest_efficiency	0.09	B	0.045–0.135	tech gain per scaled R&D
adaptability_learning_rate	0.05	B	0.025–0.075	adaptability accumulation
legitimacy_growth_rate	0.04	B	0.02–0.06	legitimacy accumulation
obsolescence_rate*	0.012	B	0.009–0.027	base technology decay rate ω_0, scaled by the pace-obsolescence factor $(1 + \omega_g g_t)$
imitation_rate*	0	B	0–0.35	Suarez & Lanzolla 2007 (slow pace: followers catch up)
imitation_pace_damping	12	B	0–25	δ_ι: how strongly a moving frontier protects the leader; catch-up is scaled by $\text{clip}(1 - \delta_\iota g_t)$ (Eq. 20)
obsolescence_pace_gain*	0	B	0–30	Henderson & Clark 1990; Tushman & Anderson 1986 (fast pace: obsolescence)
qw_technology	0.4	A	0.25–0.55	quality weight on technology
qw_knowledge	0.3	A	0.15–0.45	quality weight on knowledge
qw_adaptability	0.1	A	0–0.25	quality weight on adaptability
qw_legitimacy	0.2	A	0.1–0.35	quality weight on legitimacy
gestation_friction_penalty	0.15	A	0.05–0.3	route-symmetric temporary demand penalty fr_i while a ramp, integration, or coordination timer is active; only the window length differs by route
concentration_gain*	7	E	0–8	Arthur 1989 (increasing returns to adoption); Katz & Shapiro 1985
concentration_exponent_max	6	E	1–6	Arthur 1989, Thm II (cap that prevents winner-take-all lock-in)

continued on next page

Table 14, continued

Parameter	Base	Cls	Range	Grounding / note
`min_efficient_scale`	1	E	0.3–2.5	Klepper 1996; Klepper & Simons 1997 (scale/R&D-cost shakeout)
`scale_economy_exponent`*	0.9	E	0–1.2	Klepper 1996 (unit-cost scale economies; inert for the shakeout in the knockout panel)
`fixed_cost_scale`	0.18	A	0.09–0.27	fixed-cost coefficient ($\times\sqrt{C_i}$)
`variable_cost_scale`	0.12	A	0.06–0.18	variable-cost coefficient ($\times$ revenue)
`base_margin`	0.28	A	0.14–0.42	base gross margin
`margin_quality_weight`	0.16	A	0.08–0.24	margin sensitivity to relative quality
`margin_floor`	0.08	A	0.04–0.12	lower clamp on margin
`margin_cap`	0.45	A	0.35–0.6	upper clamp on margin
`external_finance_scale`	4.5	B	2–7	external financing amount scale
`external_finance_probability_scale`	1.8	B	0.9–2.7	financing-probability sigmoid steepness
`capacity_adjust_rate`	0.28	B	0.14–0.42	partial-adjustment speed toward target capacity
`minimum_capacity`	0.12	N	n/a	floor on production capacity
`allocation_rounds`	6	N	n/a	iterative demand-allocation rounds
`demand_overflow_fraction`	0.25	N	n/a	scaling of the final backorder overflow pass
`parent_support_buffer`	8	B	4–12	cash a sponsor keeps before supporting a unit
`parent_support_target_cash`	4	B	2–6	cash level the sponsor tops a unit up to
`parent_support_share`	0.35	B	0.15–0.55	share of a unit's loss covered as support
`exit_loss_streak_threshold`	3	B	2–5	Jovanovic 1982
`exit_share_floor`*	0.01	B	0.001–0.01	no-traction market-share floor
`exit_finance_floor`	−1.5	B	−3 to −0.5	cash floor for finance-distress exit
`feasibility_symmetry`	True	S	n/a	symmetry principle (Sec. 3.4)
`knockout_freeze_pool`	False	S	n/a	variant of K1 that freezes the pool; not used in the reported results
`knockout_no_depletion`	False	S	n/a	knockout K1, clean (Sec. 3.7)
`knockout_hold_barrier`	False	S	n/a	knockout K2, freezing the barrier; not used in the reported results
`knockout_randomize_timing`	False	S	n/a	knockout K3 (Sec. 3.7)
`knockout_homogenize_resources`	False	S	n/a	knockout K4 (Sec. 3.7)
`knockout_randomize_entry_order`	False	S	n/a	knockout K5 (Sec. 3.7)
`entry_wave_cutoff_year`	12	S	n/a	K5 late-wave cutoff year

Appendix C: Design–evidence map

Table 15 collects every major design choice in the model and places it next to the academic evidence or industry practice that motivates it. The left column states what the model does and where; the right column states why, with the supporting literature. The table operates at the level of mechanisms and functional forms; the coefficient-level anchors, governed ranges, and governance classes for every individual parameter are in Table 14.

Table 15: Design–evidence map: each model design choice (left) and the academic or industry evidence behind it (right).

Model design choice	Evidence and rationale
Market environment	

Table 15 (continued)

Model design choice	Evidence and rationale
Two coupled markets: a declining legacy arena and a growing emerging arena (Sec. 3.1)	Incumbents under disruption manage co-evolving old and new regimes simultaneously rather than facing a clean switch (Eggers and Park, 2018; Verhoef et al., 2021).
Annual time step; sixty-year horizon (Sec. 3.1)	Entry, acquisition, and alliance decisions follow roughly annual planning cycles; the horizon must contain the growth, peak, and shakeout of a full industry life cycle (Klepper, 1996).
Emerging share evolves by competitive substitution: a replicator rule on a signed fitness gap (Eq. 3)	Technological substitution follows fitness-driven share dynamics (Fisher and Pry, 1971). The signed gap lets a transition stall or fail, making disruption a treatment rather than an assumption; an automatic diffusion curve would presuppose the disruption the model exists to study.
Technology frontiers are population-derived (the maximum over active firms), not imposed (Eq. 3)	Keeps the opportunity structure endogenous to firm behavior, per the generative standard that headline conditions be produced rather than installed (Epstein, 1999).
Dominant-design strength rises with emerging share, frontier, and concentration (Eq. 4)	A regime settles once it has scale, technical coherence, and consolidation (Anderson and Tushman, 1990; Suarez and Utterback, 1995).
Entry barrier rises with installed base, quality, and concentration (Eq. 4)	Increasing returns and lock-in raise the hurdle later entrants must clear (Arthur, 1989).
One increasing-returns primitive: demand allocated by a sharpened position pos^{β}, β rising with dominant-design strength, capped (Eqs. 7, 8)	Self-reinforcing advantage (Arthur, 1989) drives both rising concentration and the shakeout, with the exit floor turning falling share into exit (Klepper, 1996); the cap reflects the observed oligopoly outcome rather than winner-take-all. Verified by knockout (Sec. 4.8).
Position formula $\text{pos}_i = Q_i\sqrt{C_i}$ with quality blending the four capability stocks (Eq. 7)	Survival requires both a competitive product and the capacity to serve demand (Klepper, 1996); the concave scale term encodes diminishing returns to size (Penrose, 1959).
Capacity-constrained allocation rounds and a sub-scale exit floor (Eq. 8; Sec. 3.5)	The industry life cycle shows a post-peak shakeout in the number of firms (Klepper, 1996); in the model it comes from the sharpened allocation together with the exit floor (Section 4.8).
De novo entry by a stage-modulated Poisson process that tapers as the design consolidates (Eq. 9)	The window-of-opportunity entry pattern of industry evolution (Suarez et al., 2015; Klepper, 1996).
Firms and initialization	
Three actor classes: legacy incumbents, emerging independents, and affiliates (Sec. 3.2)	Independents are simultaneously entrants, acquisition targets, and alliance partners, the counterparties that make the partner and acquire routes possible (Capron and Mitchell, 2012).
Four capability stocks: knowledge, technology, adaptability, legitimacy (Sec. 3.2)	The minimum set the theory requires: what a firm can do (K, T), how readily it changes (A) (Hannan and Freeman, 1984), and whether the market accepts it (L) (Stinchcombe, 1965).
Continuous resources within overlapping seed size classes, drawn from one latent maturity factor with evidence-signed loadings; rules use realized resources, not class labels (Eq. 1)	Competitive position is a stock built over time (Penrose, 1959; Dierickx and Cool, 1989); adaptability falls with age per structural inertia (Hannan and Freeman, 1984). Prevents stratification from being assumed rather than found (Epstein, 1999).

Table 15 (continued)

Model design choice	Evidence and rationale
Size enters decisions only through cash (affordability and cost-relative-to-cash); the class label is never read by any rule (Sec. 3.2)	Lets resource sorting emerge from the mechanism rather than from a label; the descriptive quantile is used only for reporting and as a control. Verified by the size-structure checks (Sec. 4.8).
Sensing	
Attention is a sigmoid of evidence-weighted signals: performance gap, share loss, frontier gap, and disruption flag raise it; customer dependence, age, and low adaptability lower it (Eq. 10)	Problemistic search intensifies under adverse performance feedback (Cyert and March, 1963; Greve, 2003); attention is biased toward existing customers and routines (Christensen and Bower, 1996; Gilbert, 2005; Tripsas and Gavetti, 2000). The sigmoid bounds attention and yields threshold-like activation.
Attention gates whether a firm evaluates entry, never which route it takes (Sec. 3.4)	Separates who searches from what searchers choose, the identification SQ1 requires; merging the two would conflate sensing with seizing.
Inverted-U slack effect in search pressure (Eq. 11)	The documented inverted-U relationship between slack and search or innovation (Bourgeois, 1981; Nohria and Gulati, 1996).
Screening	
One symmetric fit function for partners and targets: similarity, complementarity, and candidate quality (Eq. 12)	Interorganizational learning needs both overlap and complement, per relative absorptive capacity (Lane and Lubatkin, 1998); a single rule prevents either route from gaining a hidden screening advantage.
Bilateral alliance acceptance that falls with a large size gap (Eq. 13)	The appropriation-hazard and leakage logic of the relational view (Dyer and Singh, 1998).
Seizing	
Softmax (random-utility) route choice at temperature τ (Eq. 17)	The standard formal representation of probabilistic discrete choice (McFadden, 1974), matching the behavioral premise that firms satisfice under uncertainty (Cyert and March, 1963; Nelson and Winter, 1982). Best response would assume known payoffs on an unresolved frontier; reinforcement learning adds parameters with no observable counterpart.
All routes scored on one shared formula; the routes differ only in the starting position each delivers (Eq. 14)	The resource-pathways trade-off itself: build inherits capability but starts small and ramps slowly; acquisition buys a full operating position less an integration loss that grows with capability distance; partnership combines sponsors at the equity stake with coordination cost (Capron and Mitchell, 2012, 2009; Ahuja and Katila, 2001; Puranam et al., 2006; Kale and Singh, 2007).
Real-options discount on the irreversible routes only; option value on the flexible routes (Eq. 14)	Irreversible commitment forfeits the value of waiting until uncertainty resolves (Folta, 1998; Dixit and Pindyck, 1994; Kogut, 1991).

Table 15 (continued)

Model design choice	Evidence and rationale
Contractibility acts on alliance acceptance and on a premium for owning a tacit capability, not on the partner route's value (Eqs. 13, 14)	Codified capabilities can be traded at arm's length while tacit ones must be owned (Williamson, 1985; Capron and Mitchell, 2012). Firms whose strategy emphasizes value appropriation rather than value creation prefer acquisitions to alliances (Mallapragada et al., 2025); reading that finding as one of contractibility is our interpretation. Acting on acceptance avoids the choose-then-fail bottleneck that a boost to the partner route's value produced.
Symmetric feasibility threshold; a route is feasible only if affordable, with acquisition price scaling with target size (Eqs. 15, 16)	The no-privileged-route principle; acquisition pricing practice ties the price to the scale of the target (Haspeslagh and Jemison, 1991).
State-dependent route costs: build cost rises with the entry barrier, partner cost with uncertainty, acquisition cost with the target's captured demand and capacity and, in robustness runs, with dominant-design strength (Eq. 15; Sec. 4.8)	Acquisition premiums in practice scale with target value (Haspeslagh and Jemison, 1991); the friction stands in for merger review and target scarcity as an industry-practice counterfactual.
Transformation	
Learning gated by absorptive capacity, with gap-closing increments $e \cdot \mathrm{rd} \cdot ac \cdot (1 - X)$ (Eq. 20)	A firm learns from investment, partners, and targets only to the degree its related knowledge allows (Cohen and Levinthal, 1990; Lane and Lubatkin, 1998; Zollo and Winter, 2002); the concave gap-closing form is the standard asset-accumulation specification (Dierickx and Cool, 1989).
Route-specific, temporary, equal-sized frictions: build ramp, integration window, alliance coordination drag (Sec. 3.4)	Each friction is documented: internal builds ramp slowly (Capron and Mitchell, 2012), acquisitions need integration (Haspeslagh and Jemison, 1991; Puranam et al., 2006), alliances carry coordination cost (Kale and Singh, 2007); equal size enforces the symmetry principle.
Imitation catch-up and pace-obsolescence channels, off at baseline and activated only for the timing experiment and the first-mover arm of the knockout panel, $a = 3.4$ (Sec. 3.5)	Slow frontiers let followers catch up; fast frontiers obsolesce a pioneer's lead (Suarez and Lanzolla, 2007; Lieberman and Montgomery, 1998).
Parent support buffers and route-specific patience windows for young affiliates (Sec. 3.5)	The liability of newness: organizations face heightened mortality early in life and need protection while ramping (Stinchcombe, 1965; Jovanovic, 1982).
Margin band, financing, and capacity-adjustment rules (Eq. 19)	Accounting normalizations (classes A and N in Table 14), fixed to convenient scales rather than asserted as behavioral claims.
Selection	
Entity- and route-specific exit on combined distress and lack of traction; legacy parents fail on their own core-business distress (Sec. 3.5)	Noisy selection of viable firms (Jovanovic, 1982) and ecological mortality (Hannan and Freeman, 1984); a uniform rule would either remove ventures before they could ramp or shield parents whose core had failed.

Table 15 (continued)

Model design choice	Evidence and rationale
Leader and survivor thresholds applied identically across routes (Sec. 3.6); alliance value weighted by an explicit equity stake θ (Sec. 3.4)	The symmetry principle: an alliance-heavy firm must be able to lead; the explicit stake replaces an arbitrary one-half discount.
Identification and experiments	
Each generative, confound, and decomposition headline paired with a manipulation of its mechanism; the entrant ranking examined by a build-scale sweep (Secs. 3.7, 4.7, 4.8)	The generative standard: a pattern counts as grown only if it collapses or falls sharply, persists, or decomposes as its type requires (Epstein, 1999; Davis et al., 2007).
Variance-based Sobol sensitivity over the treatment parameters (Sec. 3.7)	One-factor-at-a-time analysis fails under the nonlinear interactions of agent-based models (Saltelli et al., 2010; ten Broeke et al., 2016).
Common random numbers across conditions; outcomes reported as distributions with confidence intervals (Sec. 3.7)	Rigor guidelines for agent-based research (Rand and Rust, 2011).
ODD+D documentation and seeded, hash-stamped reproducibility (App. A)	The reporting standard for agent-based models with human decision-making (Grimm et al., 2010; Müller et al., 2013).

Appendix D: Empirical anchoring

This appendix places the model's mechanisms beside three well-documented technology transitions. It is an illustrative correspondence, not an empirical test: the model is calibrated to stylized facts and uses no firm-level data, so these cases are meant to show that its primitives (route costs, the contractibility of the needed capability, the liquidity of the target pool, the pace of the frontier, first-mover timing, target-pool depletion, and increasing returns) correspond to recognizable episodes, not to estimate or validate the model against them. The three cases are chosen to span the contractibility and target-liquidity range that the model varies, from low (tacit) to high (codified) capability and from thin to deep target pools; Appendix C already ties each mechanism to its supporting literature. Table 16 summarizes the mapping, and the paragraphs below read across its rows. Dates and deal values are cited to company announcements, filings, and the business press; they are given for orientation, and the correspondence between each case and the model is our reading.

Film to digital imaging. This transition pairs a failed incumbent with a survivor. Kodak built the enabling technology early, its engineer Steven Sasson building the first self-contained digital camera in 1975 (Estrin, 2015), yet the firm largely waited, protecting its film and photofinishing business until it filed for Chapter 11 in January 2012 (Eastman Kodak Company, 2012). Fujifilm, facing the same collapse in film demand, reconfigured instead. It redeployed its film-chemistry and materials know-how through internal build, launching the Astalift cosmetics line in 2007 (FUJIFILM Holdings Corporation, 2024), and through acquisition, taking 66% of Toyama Chemical in March 2008 for about $0.98 billion to enter pharmaceuticals (FUJIFILM Holdings Corporation, 2008); by the fiscal year ended March 2011 its imaging segment was about 15% of revenue (FUJIFILM Holdings Corporation, 2011). In the model's terms the needed digital competence was of

low contractibility, and the pool of targets that could confer it was thin, so the routes that worked were internal rebuilding and diversifying acquisition. The frontier moved quickly once digital cameras took hold, and Kodak's early position in the underlying technology did not turn into a lead, a pattern that the attention gate of Eq. 10 is built to represent, since attention anchored to the legacy business delays the response, as Tripsas and Gavetti (2000) document for Polaroid (see also Christensen, 1997).

Internal combustion to electric vehicles. This is the one transition in which all four sourcing modes are visible at once. Incumbents built dedicated electric platforms (Volkswagen's MEB, General Motors' Ultium, Ford's Mustang Mach-E); they partnered on the more contractible battery layer through cell joint ventures, GM with LG Chem in a venture announced in 2019 that became Ultium Cells (Wayland, 2019), Ford with SK Innovation in BlueOval SK, announced in 2021 (Alamalhodaei, 2021), and Volkswagen through its investment in and joint venture with QuantumScape (Volkswagen Group, 2020); they acquired or bought into the tacit autonomy layer, GM acquiring Cruise in 2016 (General Motors, 2016) and Ford committing $1 billion to Argo AI in 2017 (Ford Motor Company, 2017); and Toyota largely waited, bridging with hybrids. The split maps onto contractibility: batteries were codifiable enough to source through alliances and supply contracts, whereas self-driving software was tacit and resisted integration, and both autonomy bets were later wound down, Argo AI in October 2022 (Korosec, 2022) and GM's funding of Cruise's robotaxi work in December 2024 (General Motors, 2024). Only a handful of credible autonomy and battery targets existed, so target liquidity was moderate. Tesla's first-mover advantage in battery-electric vehicles largely held (Perkins and Murmann, 2018) while the incumbents' autonomy investments faltered, a mixed outcome; in the model the early-entry gap stays positive at every pace tested and narrows only for self-timed entry at the fastest paces.

On-premise software to cloud and SaaS. This case shows the acquire-heavy pattern the model produces when targets are abundant and acquisition is cheap relative to building. Codebases and interfaces are codifiable and integrate at relatively low cost, and once SaaS emerged as a new computing platform (Cusumano, 2010; Benlian and Hess, 2011) the venture-funded sector supplied a deep, liquid pool of targets. Incumbents both built and acquired. Microsoft made Azure generally available in 2010 (Microsoft, 2010) while acquiring LinkedIn for $26.2 billion in 2016 (Microsoft, 2016) and GitHub for $7.5 billion in 2018 (Microsoft, 2018). Salesforce, itself a cloud-native entrant, bought MuleSoft for about $6.5 billion in 2018 (Salesforce, 2018), Tableau for $15.7 billion in 2019 (Salesforce, 2019), and Slack, a $27.7 billion deal announced in December 2020 and closed in July 2021 (Salesforce and Slack Technologies, 2020; Salesforce, 2021). Oracle bought NetSuite for about $9.3 billion in 2016 (Oracle, 2016); SAP bought Concur for about $8.3 billion in 2014 (SAP SE, 2014) and agreed in 2018 to buy Qualtrics for $8 billion, closing in 2019 (SAP SE, 2018, 2019); and Adobe moved Creative Suite to Creative Cloud subscriptions in 2013 (Lardinois, 2013) and acquired Magento and Marketo in 2018 (Adobe Inc., 2019). High target liquidity did not make acquisition frictionless: Adobe and Figma terminated their merger agreement in December 2023, citing no clear path to regulatory approval in Europe and the United Kingdom (Adobe Inc., 2023), the review constraint the model represents as a rising acquisition price.

Appendix E: Supplementary figures

This appendix collects a table and three figures that support the main results but are not central to the argument: the phase-map area shares under the two scoring rules (Table 17), the Sobol sensitivity indices behind the phase map (Figure 8), the composition of the emerging market by

Table 16: Illustrative mapping of the model's primitives onto three technology transitions: the routes observed and the model conditions each case resembles. The correspondence is our qualitative reading and is not an empirical test; sources are cited in the text.

Transition	Routes observed				Model conditions		
	Build	Partner	Acquire	Wait	Contract-ibility	Target liquid-ity	Pace
Film to digital imaging	Kodak (early digital camera); Fujifilm (Astalift)	–	Fujifilm (Toyama Chemical)	Kodak (protected film business)	Low	Low	Fast
Internal combustion to electric vehicles	VW MEB, GM Ultium, Ford Mach-E	Battery-cell joint ventures (GM–LG Chem, Ford–SK Innovation, VW–QuantumScape)	GM (Cruise); Ford (Argo AI, investment)	Toyota (hybrids first)	Moderate	Moderate	Fast
On-premise software to cloud and SaaS	Microsoft Azure; Adobe Creative Cloud	Independent software vendor ecosystems	Microsoft, Salesforce, Oracle, SAP, Adobe (multiple deals)	–	High	High	Fast

In the model, high contractibility favors partnering; the cloud row is acquire-heavy because of deep target liquidity and low acquisition cost.

condition (Figure 9), and the robustness of the contingency result to cost-neutral and acquisition-friction manipulations (Figure 10).

Table 17: Phase-map area shares under two scoring rules: the most leaders per cell (the map of Figure 5) and the most leaders per entrant, leader share$_r$/route share$_r$. Rerun of the four 12×12 grids with the same three seeds and fifty-year horizon; the most-leaders shares reproduce the stored maps exactly. Brackets give bootstrap 95% intervals from 1,000 resamples of the three replicate seeds; with only three seeds the resampled argmax can take few distinct values, so a point estimate can sit on an interval bound. Agreement is the fraction of cells where the two rules pick the same route.

Plane	Most leaders			Leaders per entrant			
	Build	Partner	Acquire	Build	Partner	Acquire	Agreement
Build cost × partner cost	0.18	0.16	0.66	0.03	0.16	0.81	0.56
	[0.16, 0.22]	[0.13, 0.22]	[0.58, 0.69]	[0.01, 0.08]	[0.16, 0.25]	[0.69, 0.81]	
Contractibility × build cost	0.19	0.03	0.78	0.01	0.08	0.92	0.74
	[0.16, 0.27]	[0.01, 0.12]	[0.60, 0.83]	[0.00, 0.03]	[0.08, 0.25]	[0.72, 0.92]	
Build cost × pace	0.18	0.00	0.82	0.01	0.01	0.98	0.81
	[0.16, 0.25]	[0.00, 0.03]	[0.72, 0.83]	[0.01, 0.06]	[0.01, 0.12]	[0.81, 0.98]	
Target liquidity × uncertainty	0.00	0.00	1.00	0.00	0.01	0.99	0.99
	[0.00, 0.03]	[0.00, 0.14]	[0.83, 1.00]	[0.00, 0.04]	[0.01, 0.16]	[0.80, 0.99]	

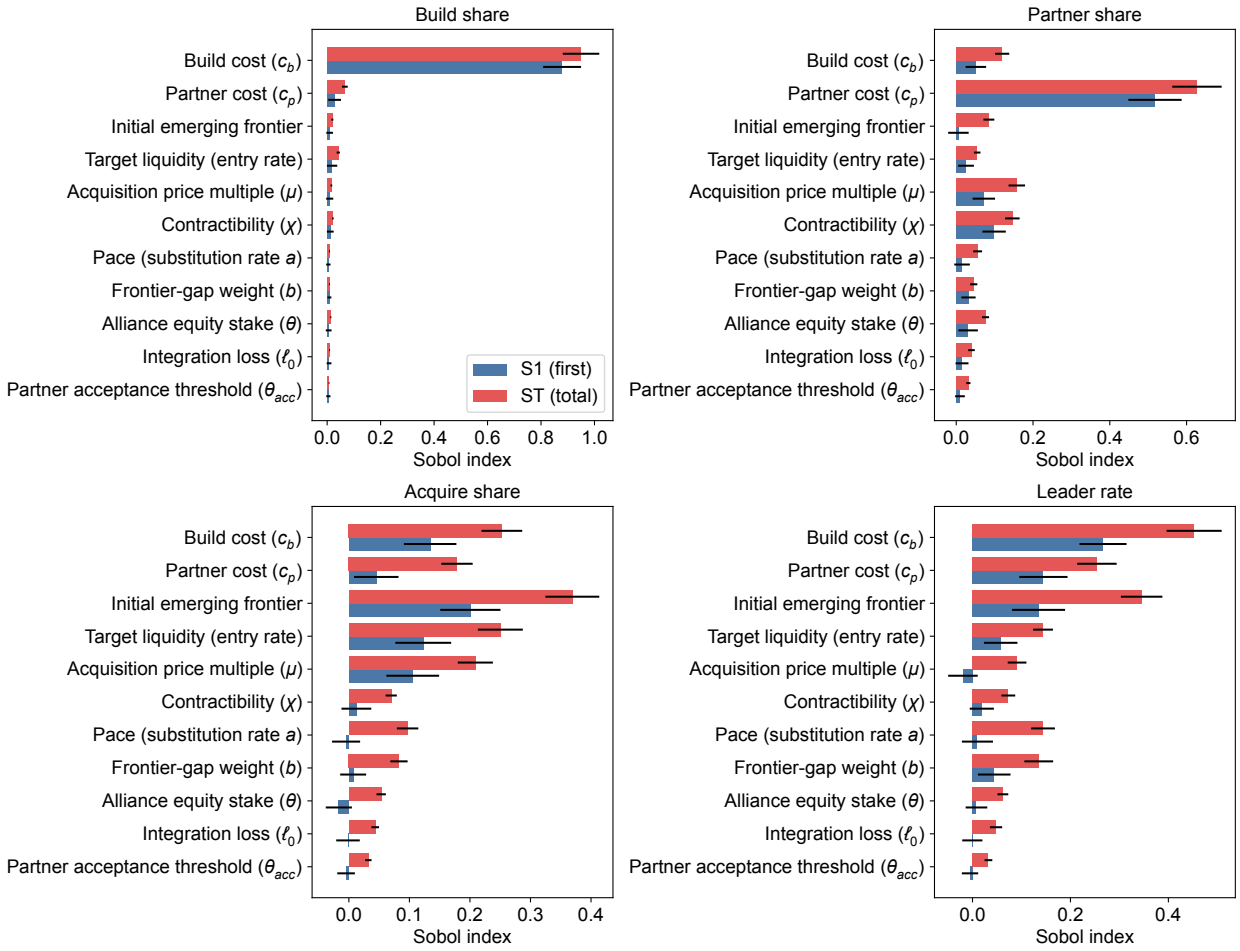


Figure 8: Sobol first-order (S1) and total-order (ST) sensitivity indices per outcome, over the eleven treatment parameters, with bootstrap confidence intervals. The build share is governed almost entirely by the build cost and the partner share mainly by the partner cost; the acquire share is the exception, governed by the initial emerging frontier rather than by the route costs, because the acquisition fixed cost is held at its low baseline in this design.

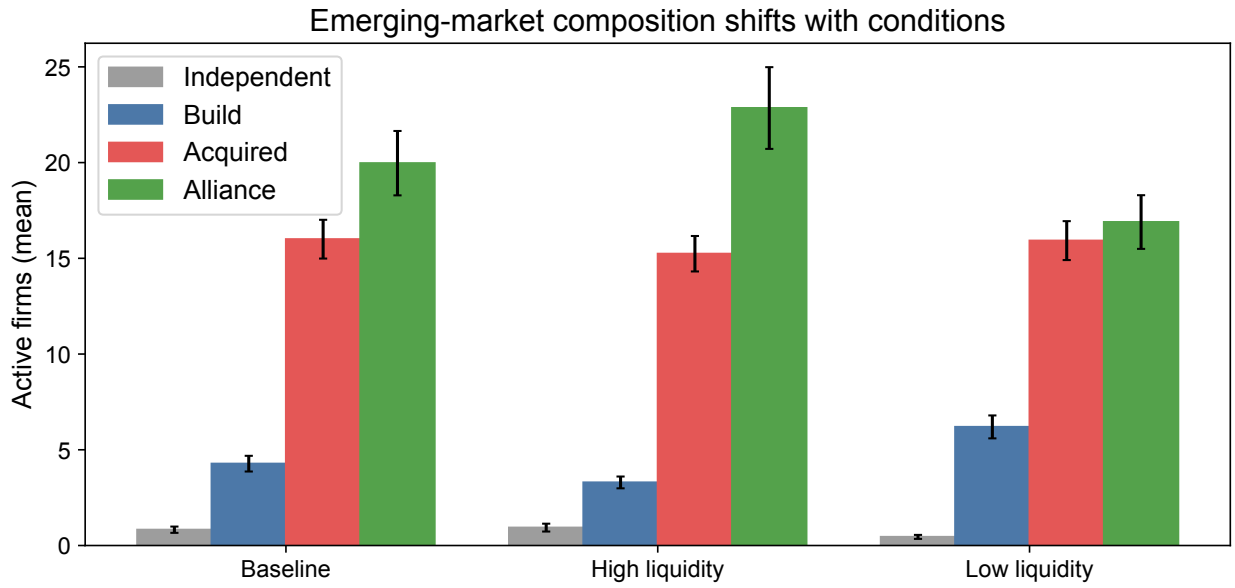


Figure 9: Composition of the emerging market by condition: the mean number of active independents, internal builds, acquired units, and alliance units at the end of the run, at baseline and at low and high target liquidity (200 runs per condition; error bars are 95% confidence intervals across runs). The affiliate-persistence rates quoted in Section 4.6 are computed from the same runs.

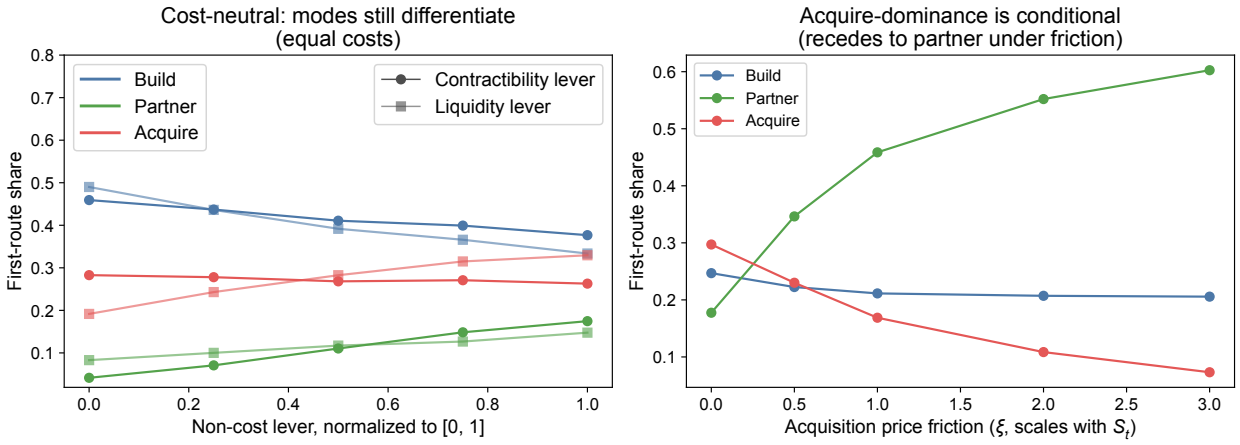


Figure 10: Robustness of the contingency result (first-route shares over all incumbents, means across 150 runs per point). Left: with build and partner costs held equal, the mode mix still responds to the non-cost levers of contractibility and target liquidity, so the route mix is not a cost artifact. Right: as the acquisition-price friction that rises with dominant-design strength is increased, acquisition recedes and partner takes over, so acquisition's prominence is conditional rather than structural.