WORKING PAPER · REVISED EDITION

Epistemological Velocity Theory: Frontier Advantage under Adversarial Formalization

A theory of strategy for firms whose formalizable knowledge is being operationalized by machines faster than they can renew it.

AUTHOR
Brad Heitmann

AFFILIATION
Salt Lake City, Utah

VERSION
Revised working paper, August 2026

SUPERSEDES
Draft of June 2025

00 ABSTRACT

The theory-based view holds that firms create value by forming novel theories that reveal opportunities invisible from data alone. This paper extends that view with two elements it does not contain: a rate variable, and an adversary. The rate variable is epistemological velocity, ω, the frequency with which a firm revises the framework through which it reads its environment. The adversary is formalization, an external process, now dominated by machine learning systems, that converts any problem derivable from existing frameworks into a procedure any actor can run. The paper partitions the firm's problem space into a formalized set K, a derivable frontier D, and a non-derivable frontier N, and specifies the flows between them. Advantage created on D decays with the formalization half-life ln 2 / φ; advantage created on N persists until the firm's own conjecture actualizes it. The share of new frontier that is non-derivable, λ, is treated as an empirical and declining quantity rather than a metaphysical constant, and the condition under which the theory fails is stated. The rate of frontier-advantage creation is $S = \omega \cdot Q \cdot \lambda_{eff} - c(\omega)$, where Q weights revisions by their distance from the prior framework and c is a convex coherence cost; velocity therefore has an interior optimum, not a maximum. Sourced evidence on knowledge half-lives and on the measured growth rate of machine task horizons anchors the two exogenous parameters. A measurement program is proposed. The empirical figures in this edition cite primary sources, and the model figures state their parameters.

§ CONTENTS

01 The claim, and what it adds

02 Prior work the theory stands on

03 The model

04 Evidence on the exogenous parameters

05 A process-philosophical reading

06 Measurement

07 Organizational implications

08 Objections

09 Conclusion

A Claims ledger

B Errata against the June 2025 draft

R References

01 THE CLAIM, AND WHAT IT ADDS

The claim, and what it adds

Strategy theory already has an account of how firms see what data cannot show. It does not have an account of how fast that seeing has to be renewed when a machine is converting the last insight into a commodity.

Felin and Zenger (2017) describe economic actors as theorists: a firm's value comes from a theory that picks out problems and assets the market has not yet priced, and the theory is tested by the firm's own experiments rather than read off the data. Zenger (2013) gives the practitioner form. Felin and Holweg (2024) then argue that present machine learning systems cannot do this work, because they are bound to backward-looking statistical association while human cognition can hold beliefs ahead of the evidence and act on them. Taken together these papers describe a static asymmetry. The firm theorizes; the machine predicts within what has been theorized.

This paper accepts the first paper as its foundation and treats the second as a hypothesis about a parameter rather than a settled premise. It adds two elements.

The first is a rate. A theory of value does not hold its edge; it is copied, priced in, or overtaken. What matters strategically is not the possession of a theory but the frequency with which a firm replaces it with a better one, weighted by how far each replacement moves the firm's opportunity set. Call that epistemological velocity, ω, measured in framework revisions per year and weighted by a quality factor Q.

The second is an adversary. Once a problem can be stated within an existing framework, it can be proceduralized, and once proceduralized it can be run by any actor with access to the procedure. Machine learning systems have made the proceduralization step fast and cheap across a widening class of problems. Call that formalization, with rate φ. Formalization does not attack the firm; it attacks the durability of anything the firm has already made explicit.

The interaction of these two elements is the theory. A firm's frontier advantage is created by theory revision and destroyed by formalization. The share of revision that lands beyond formalization's reach, λ, is the human edge, and the paper's central methodological commitment is that λ is a measurable and declining quantity, not a property of consciousness secured by argument. The theory says what happens at every value of λ, including zero.

Three things follow that the theory-based view does not say on its own. Velocity has an interior optimum, because revision has a coherence cost that rises faster than its yield. Advantage built on derivable problems has a half-life set by an external parameter the firm does not control. And the strategic question shifts from what is our theory to how fast must our theory move, and where must it move to, given what can now be formalized.

02 PRIOR WORK THE THEORY STANDS ON

Prior work the theory stands on

The June 2025 draft of this paper claimed that strategic thought had privileged being over becoming. That claim is withdrawn. The literatures below already treat firms as processes and advantage as a rate; the contribution here is narrower and is stated against them.

The theory-based view

Felin and Zenger (2017) hold that value creation begins with a theory: a causal conjecture about how assets, problems and activities could be combined, held before the evidence for it exists. The theory directs search and experimentation and makes some opportunities visible to the firm that hold it and invisible to firms that do not. This paper's concept of a firm's epistemological state Ψ is exactly the firm's theory of value in that sense. Where this paper departs is in treating the theory as perishable on a measurable clock, and in modeling a competitor that is not another theorist but a formalizer.

Data-based and theory-based cognition

Felin and Holweg (2024) argue that large language models and related systems reason from data and cannot generate forward-looking causal beliefs that outrun their training distribution, while humans can and do. This paper uses that argument as the source of the human-edge hypothesis and then relocates it. Rather than an axiom, the asymmetry becomes a parameter $\lambda(t)$ whose value is to be estimated and whose trajectory is to be watched. Section 8 states what the theory implies if the asymmetry closes.

Process organization studies

Tsoukas and Chia (2002) reconceived organizational change as the normal condition of organizing, drawing on Bergson and James. Hernes (2014) built an account of organization on Whitehead's process metaphysics. Nonaka, Toyama and Hirata (2008) offered a process theory of the knowledge-based firm. MacKay, Chia and Nair (2021) developed a process-philosophical approach to strategy itself. The application of process philosophy to organizations and to strategy is therefore established, and the present paper does not claim it. Section 5 uses Whitehead and Bergson as an interpretive vocabulary for a model that stands without them.

Dynamic capabilities and absorptive capacity

Teece, Pisano and Shuen (1997) located advantage in the capacity to integrate, build and reconfigure competences; Eisenhardt and Martin (2000) characterized that capacity as identifiable routines whose value depends on market velocity. Cohen and Levinthal (1990) defined absorptive capacity as the ability to recognize, assimilate and apply new external knowledge. Both are rate theories. This paper isolates one kind of reconfiguration, revision of the framework itself as opposed to improvement within it, and introduces a specific external process that erodes the value of everything the framework has already made explicit. March (1991) supplies the trade-off that makes velocity an optimization rather than a maximization.

Uncertainty and the un-prestatable

Knight (1921) separated risk, which can be measured, from uncertainty, which cannot. Popper (1957) argued that the future growth of knowledge cannot be predicted by rational or scientific means, since predicting it would amount to possessing it. Kauffman (2000) named the adjacent possible and argued that the space of what can come next cannot be prestated. Felin, Kauffman, Koppl and Longo (2014) brought that argument into strategy, against the fitness-landscape picture of search over a given space. These are the formal reasons for believing that a non-derivable frontier exists at all, and this paper's N is that frontier under a different name.

Learning curves and knowledge decay

Argote (2013) documents learning within a framework: performance improves with cumulative experience and depreciates when experience is lost. Arbesman (2012) popularized the observation that facts decay on measurable schedules, drawing on Burton and Kebler (1960), Poynard and colleagues (2002) and related work. Section 4 uses those primary studies directly. Kuhn (1962) supplies the distinction between work within a paradigm and replacement of it; his term is retained for the distinction and dropped as a unit of measure.

03 THE MODEL

The model

State space

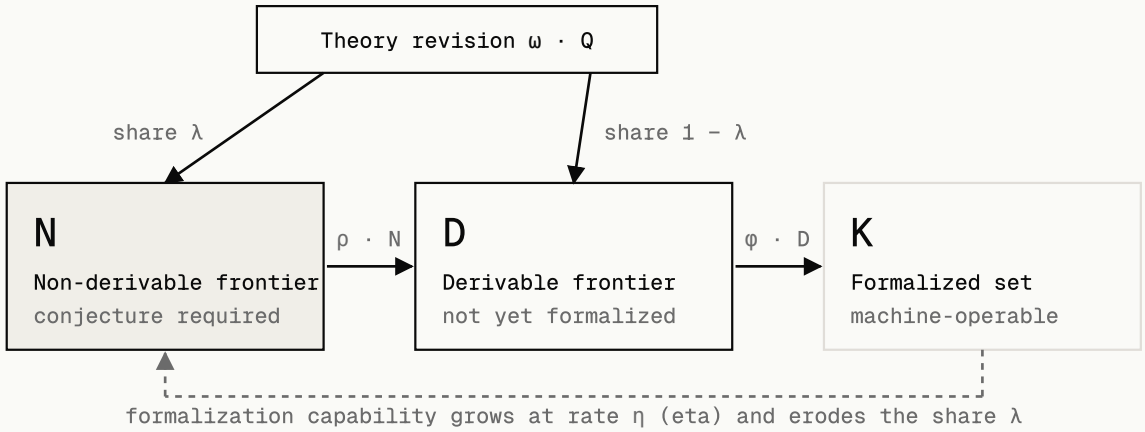
Let the value-relevant problem space of a firm be partitioned at each moment into three sets.

K, the formalized set: problems for which a procedure exists that any actor with access to it can execute. Machine learning systems operate here. Advantage on K is contestable at the cost of the procedure.

D, the derivable frontier: problems not yet formalized but statable within existing frameworks, so that formalization can reach them without new conjecture. This is the frontier in the first of two senses the June 2025 draft ran together: the not-yet-known.

N, the non-derivable frontier: problems that cannot be stated within existing frameworks and require a conjecture that is not derivable from them. This is the frontier in the second sense: Kauffman's un-prestatable adjacent possible, Whitehead's creative advance. Only theory revision reaches it.

FIGURE 1 · THE STATE SPACE AND ITS FLOWS



Firm advantage lives in N and D; K is contestable by any actor with formalization capacity. Derivable advantage decays with half-life ln 2 / φ; non-derivable advantage persists until actualized.

Theory revision at rate ω·Q creates new frontier, of which a share λ is non-derivable and lands in N while 1 − λ lands in D. Conjecture actualizes N into D at rate ρ. Formalization moves D into K at rate φ. Growth in formalization capability, η, shrinks λ over time.

MECHANISM DIAGRAM; NO DATA.

Flows

Theory revision at velocity ω, with quality Q, generates new frontier at rate ω·Q. A share λ of that new frontier is non-derivable and enters N; the remainder enters D. The firm's own conjectural work actualizes non-derivable problems into derivable ones at rate ρ. Formalization converts derivable problems into formalized ones at rate φ. The dynamics are

$$dN/dt = \omega \cdot Q \cdot \lambda - \rho \cdot N \tag{1}$$

$$dD/dt = \rho \cdot N + \omega \cdot Q \cdot (1 - \lambda) - \varphi \cdot D \tag{2}$$

$$dK/dt = \varphi \cdot D \tag{3}$$

The steady states are $N^* = \omega \cdot Q \cdot \lambda / \rho$ and $D^* = \omega \cdot Q / \varphi$. The frontier the firm can hold, $N^* + D^*$, therefore scales with velocity and shrinks with the formalization rate; and the non-derivable part of it scales with λ. Nothing in the model normalizes the space to a fixed total. The June 2025 draft's simulation did so, and thereby could not represent the claim that human theorizing expands the space; equation (3) makes K grow without bound while (1) and (2) determine whether the frontier ahead of it grows or collapses.

Effective human edge

Advantage created on D is not worthless; it is transient. If the firm's planning horizon is T, the fraction of derivable advantage still uncontested at the horizon is $e^{-\varphi T}$, the survival function of an exponential decay with half-life ln 2 / φ. The effective share of new frontier the firm can count on is then

$$\lambda_{eff} = \lambda + (1 - \lambda) \cdot e^{-\varphi T} \tag{4}$$

As φ approaches zero, λ_{eff} approaches one and every revision counts: this is the classical world in which a good idea, once had, stays valuable for the life of the firm. As φ grows large, λ_{eff} collapses to λ and only non-derivable revision counts. Equation (4) is the precise form of the claim that machine formalization compresses the half-life of formalized advantage.

Cost and the interior optimum

Revision is not free. A framework replacement abandons partly amortized routines, forces re-coordination and spends the coherence that lets an organization act as one. Following March (1991), let that cost be convex in velocity, $c(\omega) = \kappa \cdot \omega^2$. The rate of frontier-advantage creation is

$$S(\omega) = \omega \cdot Q \cdot \lambda_{eff} - \kappa \cdot \omega^2 \tag{5}$$

with an interior optimum at $\omega^* = Q \cdot \lambda_{eff} / 2\kappa$ and a peak $S^* = (Q \cdot \lambda_{eff})^2 / 4\kappa$. Two consequences are immediate. A firm whose λ_{eff} falls, because formalization has reached its territory, should slow its revision rate, not raise it; churn without reach is pure cost. And the value of raising quality Q or protecting λ enters the peak squared, while the value of cutting coherence cost enters only linearly.

FIGURE 2 · FRONTIER-ADVANTAGE RATE AGAINST VELOCITY, THREE VALUES OF EFFECTIVE HUMAN EDGE



Equation (5) with Q = 1 and κ = 0.05. Optima at ω* = 2, 5 and 8 revisions per year for λ_{eff} = 0.2, 0.5 and 0.8, with peak rates 0.2, 1.25 and 3.2. Beyond the optimum, additional velocity destroys value.

MODEL FIGURE; PARAMETERS STATED; COMPUTED FROM EQUATION (5).

Endogenous λ

The June 2025 draft secured λ by assigning machine systems to Bergson's intellect and reserving intuition for human consciousness. That move is withdrawn. Whether a system trained on the written record of human framework changes can produce one is an empirical question, and the draft's own simulation, which made the formalization rate a slider, already conceded that the boundary moves. Let formalization capability grow at rate η, so that problems which were non-derivable relative to last year's tools become derivable relative to this year's. Then

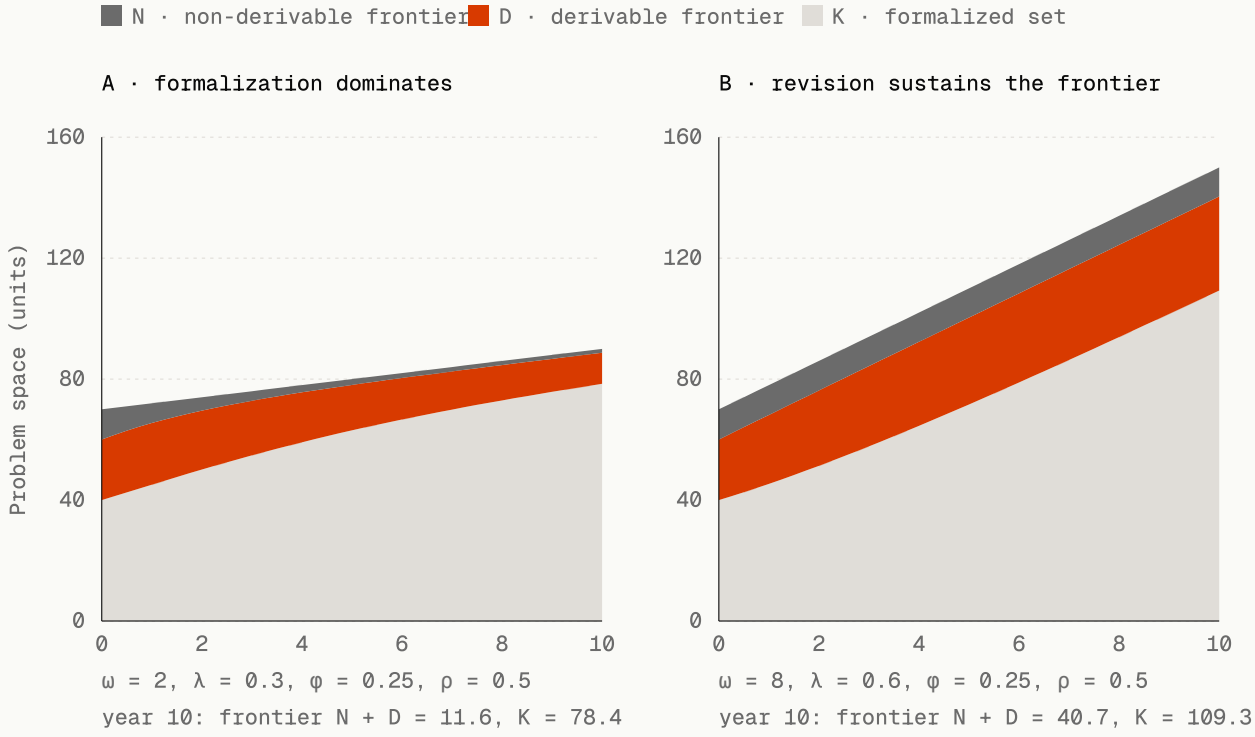
$$\lambda(t) = \lambda_\infty + (\lambda_0 - \lambda_\infty) \cdot e^{-\eta t} \tag{6}$$

where λ_0 is the present share and λ_∞ is the floor toward which it declines. The floor is the theory's central empirical parameter. If $\lambda_\infty > 0$, some class of framework revision remains beyond formalization indefinitely, and the theory-based view survives the machine age in modified form. If $\lambda_\infty = 0$, every revision is eventually derivable, S reduces to $\omega \cdot Q \cdot (1 - \lambda) \cdot e^{-\varphi T} - \kappa\omega^2$, and advantage is transient everywhere; the strategic contest becomes a race in formalization capacity, which is a resource-based account with compute as the resource.

FALSIFICATION CONDITION

The theory is falsified as a theory of durable advantage if machine systems are observed to produce framework revisions with quality Q comparable to human revisions, at comparable rates, across the classes of problem on which firms compete, so that the estimated floor λ_∞ is indistinguishable from zero. The theory would still describe the dynamics of that world, but its prescription would reduce to formalization capacity, and its name would be misleading. The observation that would confirm the theory is a measured λ(t) that declines and then stabilizes above zero.

FIGURE 3 · FRONTIER TRAJECTORIES UNDER TWO REGIMES, TEN-YEAR HORIZON



Equations (1) to (3) integrated from N = 10, D = 20, K = 40 with φ = 0.25 (formalization half-life 2.8 years) and ρ = 0.5. Regime A, ω = 2 and λ = 0.3: the frontier N + D falls from 30 to 11.6 units while K absorbs it. Regime B, ω = 8 and λ = 0.6: the frontier rises to 40.7 units and the total space grows faster. The space is not normalized; K grows in both regimes.

MODEL FIGURE; PARAMETERS STATED; EULER INTEGRATION, STEP 0.01 YEAR.

What the equation is and is not

S is the rate at which a firm creates frontier advantage. It is not strategy, and it is not total advantage. A firm with ω = 0 has S = 0 and may still hold positional, scale or execution advantages of the kind the positioning and resource-based views describe. The claim of this paper is narrower: that those advantages, insofar as they rest on formalizable knowledge, now decay at a rate set by φ, and that the only advantage a firm can renew on its own clock is the one described by S. The June 2025 draft's headline, Strategy = Velocity × Human Edge, was a category error, and its Table 1 quietly introduced a third factor absent from the headline. Equation (5) replaces both.

Units

The draft measured ω in paradigm shifts per year and plotted values of ten. A Kuhnian paradigm shift is rare, community-wide and incommensurable with its predecessor; a firm does not have ten of them in a year. The unit here is the framework revision: a change in the firm's theory of value that alters which problems it can state. Q weights each revision by the distance it moves the opportunity set, on a scale for which Section 6 proposes an operationalization.

Evidence on the exogenous parameters

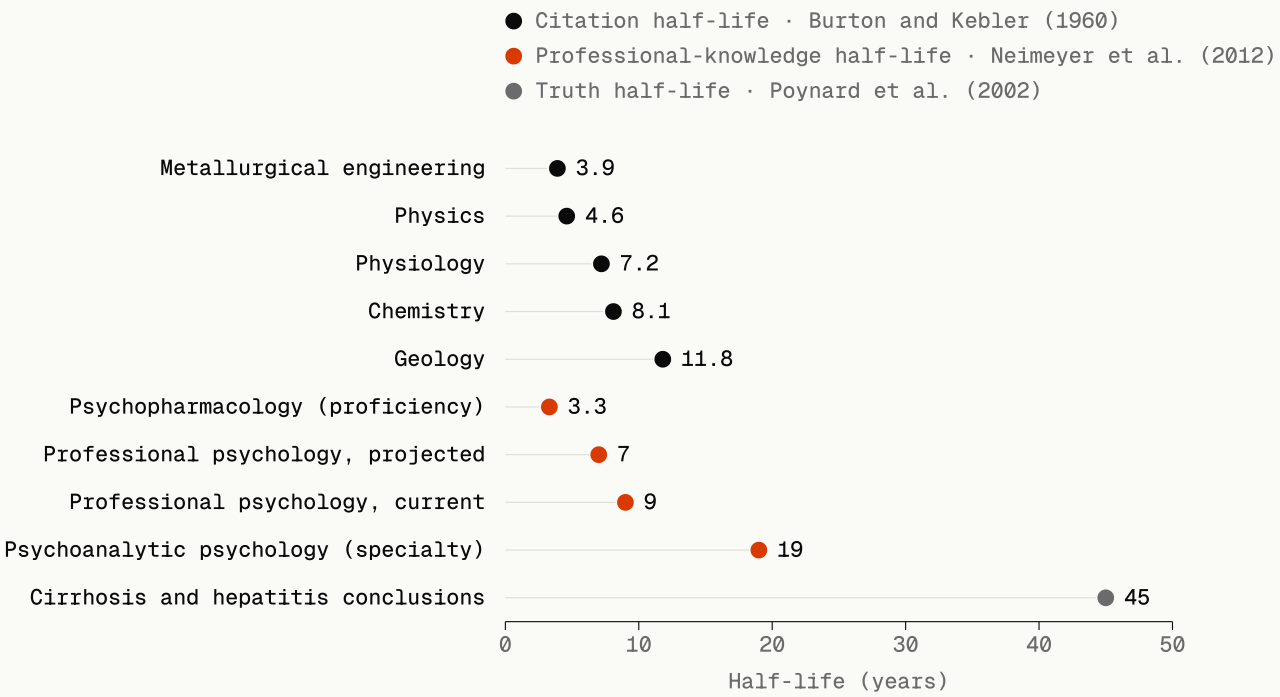
Two parameters are outside the firm's control: the decay of formalized knowledge, which sets the baseline against which φ is measured, and the rate at which formalization capability grows, η. The June 2025 draft presented four figures on these questions. Two were composed of invented series labeled as empirical evidence, one mixed sourced and unsourced points under misdated captions, and one was a share model that contradicted the theory it illustrated. All four are withdrawn. The two figures below show only values that trace to a primary source, and the constructs each source measures are named, because they are not the same construct.

Half-lives of knowledge

Three primary studies measure the decay of knowledge, each with a different construct. Burton and Kebler (1960) measured citation half-lives, the time by which half of the literature currently cited in a field was published, and reported values from 3.9 years in metallurgical engineering to 11.8 years in geology, with physics at 4.6, physiology at 7.2 and chemistry at 8.1. Neimeyer, Taylor and Rozensky (2012) polled specialty leaders in professional psychology on the half-life of knowledge in their fields and found an overall estimate of nearly nine years, expected to fall to just over seven within a decade, with a range from 3.3 years in psychopharmacology to 19 years in psychoanalytic psychology. Poynard and colleagues (2002) followed 474 conclusions from the cirrhosis and hepatitis literature and found that 60 percent were still held true in 2000, with a half-life of truth of 45 years and no survival advantage for conclusions from methodologically stronger studies.

Two things in these data bear on the model. First, the constructs differ by an order of magnitude, and the construct closest to K, what is still held true, decays slowest; what decays fastest is attention, the literature people cite. A firm's formalized knowledge does not stop being true when it is formalized; it stops being scarce. Second, the specialty-level range within one profession, 3.3 to 19 years, is wider than the range between professions, which suggests that φ is set at the level of problem class, not industry, and should be measured there.

FIGURE 4 · MEASURED HALF-LIVES OF KNOWLEDGE ACROSS THREE CONSTRUCTS



Citation half-lives (Burton and Kebler 1960, five fields as reported in secondary summaries of the paper); professional-knowledge half-lives from a Delphi poll of specialty leaders (Neimeyer, Taylor and Rozensky 2012, overall current and projected estimates and the two extremes); truth half-life of published conclusions (Poynard et al. 2002). The three constructs are not commensurable and are plotted together only to show their range.

DATA FIGURE; EVERY VALUE SOURCED; SEE THE CLAIMS LEDGER.

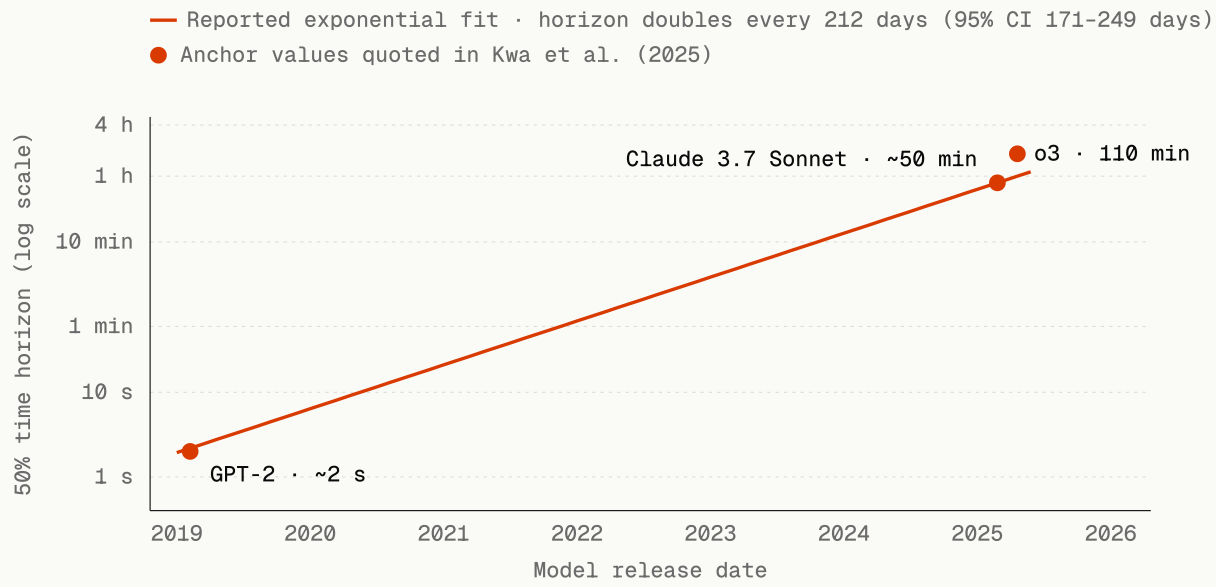
The widely repeated series that puts the half-life of an engineering degree at 35 years in 1930 and about 10 in 1960 traces to a trade-press essay and has no located primary study behind it. It appeared in the June 2025 draft with an IEEE attribution and is not used here.

The growth of formalization capability

The parameter η has no direct historical measure, because the construct is new. The nearest proxy with a published estimate is the growth of the task-completion time horizon of frontier machine learning systems. Kwa and colleagues (2025) measured, across 170 software and research tasks with human baselines, the length of task at which a system succeeds half the time, and reported that this horizon has doubled approximately every seven months since 2019, from about two seconds for GPT-2 to roughly fifty minutes for Claude 3.7 Sonnet in early 2025, with a later system at 110 minutes. The doubling time is their fitted parameter; Figure 5 redraws the fit and marks the anchor values they report.

If the proxy is even roughly right about η, the class of derivable problems is expanding at a rate that no organizational learning curve in Argote's sense approaches. That is the empirical basis for treating formalization as an adversary rather than as a tool, and for treating λ as declining. It is not evidence about λ_∞, which the proxy cannot see: a time horizon measures how long a system can work inside a stated problem, not whether it can restate the problem.

FIGURE 5 · REPORTED GROWTH OF THE 50 PERCENT TASK-COMPLETION TIME HORIZON, 2019 TO 2025



The line is the exponential fit reported by Kwa et al. (2025), a doubling time of 212 days (95 percent confidence interval 171 to 249 days in the first version of the paper), anchored at the reported value for Claude 3.7 Sonnet. Marked points are values quoted in the paper's text. The measure is a software-task horizon and is a proxy for η, not a measurement of it.

DATA FIGURE; FIT AND ANCHORS AS REPORTED BY THE SOURCE; NO VALUES RE-ESTIMATED HERE.

05 A PROCESS-PHILOSOPHICAL READING

A process-philosophical reading

The model above does not depend on any metaphysics. The June 2025 draft made Whitehead and Bergson its foundation and then extended their terms beyond what they say. This section keeps them as an interpretive vocabulary, uses only quotations that have been located, and marks each extension as one.

Whitehead

In Process and Reality, Whitehead's Category of the Ultimate is the triad of creativity, many and one, and the metaphysical principle he draws from it is that novel unities continually arise from given multiplicities:

> The ultimate metaphysical principle is the advance from disjunction to conjunction, creating a novel entity other than the entities given in disjunction. … The many become one, and are increased by one.
>
> WHITEHEAD, PROCESS AND REALITY, CORRECTED EDITION, P. 21

Read against the model, K is the objectified past, the settled data that any subsequent occasion can prehend; D is what can be reached by conformal feeling, prehension that reproduces the form of what is given; N is what requires a conceptual feeling of a possibility not yet actualized. Whitehead's propositions, lures for feeling that present how things might be rather than how they are, are the form a firm's theory of value takes before it is tested. The reading of organizations as Whiteheadian societies with personal order, patterns of inheritance that keep identity through change, follows Hernes (2014) and is not claimed as new.

One term from the draft is corrected. Whitehead's Category of Reversion concerns the phase of concrescence in which novel conceptual feelings arise as relevant alternatives to what was physically felt. The draft used conceptual reversion to mean an organizational inversion in which background assumptions become foreground variables. That is an extension by analogy, not a use of the term, and is labeled as such where it recurs in Section 7.

Bergson

Bergson's distinction between intellect, which spatializes and fragments in order to act on matter, and intuition, which grasps duration from within, is the draft's source for the human edge. His definition of intuition in the Introduction to Metaphysics is the kind of intellectual sympathy by which one places oneself within an object in order to coincide with what is unique in it. The draft quoted him as calling intuition the mind's direct participation in the universal becoming; that sentence has not been located in the standard translations and is withdrawn. The draft also closed with the line that the universe is a machine for creating novelty. The attested sentence, the last of The Two Sources of Morality and Religion, describes the essential function of the universe as a machine for the making of gods, and says something different.

What the model takes from Bergson is not the reservation of intuition to human consciousness, which Section 3 converts into a parameter, but the distinction between two temporal modes. Formalized execution runs in what he called spatialized time, homogeneous and divisible; theory revision runs in duration, in which the past interpenetrates the present and the outcome cannot be prestated. Section 7 draws an organizational-design implication from that distinction.

06 MEASUREMENT

Measurement

A theory whose thesis is that becoming can be measured cannot present illustrations as measurements. This section proposes how each construct in the model could be estimated, and Table 1 sets the indicators against the constructs. None of it has been done yet; the section is a program.

Velocity and quality

A framework revision changes which problems a firm can state, so its trace is linguistic: new categories enter the firm's documents and old ones stop doing work. Ψ can be represented as the distribution of problem-defining vocabulary in a firm's strategic documents over time, and ω as the rate of phase transitions in that distribution, detected by the same text-similarity methods Kelly, Papanikolaou, Seru and Taddy (2021) used to identify breakthrough patents by their dissimilarity to prior work and similarity to subsequent work. Applied to a firm's own patents and filings, that method yields Q directly: a revision's quality is its measured distance from the firm's prior art and the degree to which later work builds on it. Q is thereby estimated independently of market outcomes, which removes the circularity the draft had between novelty and value.

Human edge

λ is the share of a firm's new frontier that current formalization capability cannot reach. It can be estimated by classification: for each framework revision in a period, whether a state-of-the-art system given the firm's prior documents could have generated the revision's problem statement. The classification is imperfect, and its imperfection is the point; a rising rate at which the answer is yes is the direct measure of η, and the level at which the rate stops rising is the estimate of λ_∞.

Formalization rate

φ is the rate at which a firm's derivable problems become procedures others can run. Its indicators are the lag between a firm's disclosure of a method and its appearance in competitors' outputs, the share of a firm's decision loops closed without human intervention, and the decay of price premia on services the firm once performed by judgment. Because Section 4 found within-profession variance larger than between-profession variance, φ should be estimated at the level of problem class.

Cost

κ is visible in what revision destroys: the write-off of routines, the fall in coordination metrics after a reorganization, the length of the interval before a new framework reaches the productivity of the old. Argote (2013) supplies the learning-curve methods for measuring that interval.

TABLE 1 · CONSTRUCTS AND INDICATORS

CONSTRUCT	DEFINITION	ESTIMATOR	OBSERVABLE INDICATORS
Ψ, ω	Firm's theory of value; rate of its revision	Phase transitions in problem-defining vocabulary over time	New categorical frameworks; retired categories; novel problem formulations in filings and plans
Q	Distance a revision moves the opportunity set	Text dissimilarity to prior work, similarity to subsequent work (Kelly et al. 2021)	Patents requiring new classification; filings citing no prior internal art
λ, η, λ_∞	Share of new frontier beyond current formalization; its erosion; its floor	Classification of each revision as machine-reproducible from prior documents	Rate at which the classification returns yes; level at which that rate plateaus
φ	Rate at which derivable problems become executable procedures	Disclosure-to-imitation lag by problem class	Decision loops closed without human input; decay of judgment premia; skill-obsolescence curves
κ	Convex coherence cost of revision	Post-revision productivity interval (Argote 2013)	Routine write-offs; coordination-metric dips; time to prior productivity

CONSTRUCT	DEFINITION	ESTIMATOR	OBSERVABLE INDICATORS
S	Rate of frontier-advantage creation	Equation (5) from the estimates above	New-category creation; uncontested pricing on non-derivable work; strategy lifecycle length

07 ORGANIZATIONAL IMPLICATIONS

Organizational implications

Transformation curves

Learning curves in Argote's sense describe improvement within a framework and plateau as the framework's possibilities are exhausted. A transformation curve tracks the rate of framework replacement and is discontinuous: each revision resets the learning curve beneath it. The two curves compete for the same resources, and equation (5) says the split between them has an optimum that depends on λ_{eff}. A firm whose problem classes have high φ and low λ should invest in the learning curve and accept that its advantage is transient. A firm with problem classes formalization has not reached should invest in the transformation curve, up to the point where coherence cost overtakes yield. The draft's observation that high-velocity firms exhibit shorter strategy lifecycles is retained; under the model it is a consequence, not a paradox.

Two temporal modes

Bergson's distinction between spatialized time and duration corresponds to two kinds of organizational work. Execution on K and D runs on clock time and is increasingly delegated to formalized systems. Revision runs in duration, and its output cannot be scheduled. The structural implication is the one O'Reilly and Tushman (2004) drew for exploration and exploitation: the two modes need separate units, separate metrics and separate protection, because measured on execution metrics the revision unit always looks like waste. What the present model adds is the sizing rule. The revision unit's budget should track ω*, which falls as λ_{eff} falls, so a firm whose territory is being formalized should shrink and concentrate its revision effort on the problem classes where λ remains high, not spread it.

Identity through revision

A firm that revises its framework at rate ω needs something that does not revise, or it stops being one firm. Whitehead's societies with personal order supply the image: continuity kept by a pattern of inheritance rather than by fixed features. In organizational terms the invariant is not the theory of value but the standard by which theories are judged and replaced, the firm's rule for what counts as a better theory. That rule is what the draft called, by extension of a Whiteheadian term, a metamorphic structure; the extension is kept and marked. The coherence cost κ is largely the cost of lacking such a rule, since without it every revision is a contest over identity rather than over the theory.

08 OBJECTIONS

Objections

Machines will do the theorizing

This is the objection the model is built to survive. If machine systems come to produce framework revisions of human quality, $\lambda_\infty = 0$, and the theory says so and says what follows: durable advantage ceases to exist and the contest is over formalization capacity. The theory is falsified as a theory of durable advantage and confirmed as a description of the dynamics. Felin and Holweg (2024) give reasons to expect $\lambda_\infty > 0$; the position here is that those reasons describe present systems and that the parameter should be measured rather than argued.

The equation excludes most advantage

It does, by design. S is the rate of one kind of advantage creation. The claim is not that other kinds do not exist but that the ones resting on formalizable knowledge now decay on a clock the firm does not set, and that S is the only term the firm can renew on its own.

The state function is undefined

Ψ is the firm's theory of value in the sense of Felin and Zenger (2017), represented operationally as the distribution of problem-defining vocabulary in the firm's strategic documents. That is a definition the audience already grants exists. Its derivative is a rate of phase transitions in that distribution, which Section 6 makes measurable.

Novelty is defined by its payoff

Not here. Q is estimated from textual distance to prior and subsequent work, before and independently of market outcomes. Whether high-Q revisions create value is then a testable proposition rather than a definition.

This is the theory-based view in new clothes

The vocabulary is borrowed from process philosophy, and the foundation is the theory-based view; the paper says so. What is not in the theory-based view is the rate, the adversary, the half-life of derivable advantage, the interior optimum, and the falsification condition. Those are the paper.

09 CONCLUSION

Conclusion

A firm's theory of value is perishable, and the thing that perishes it has changed. It used to be imitation by other theorists, which took years. It is now formalization by systems whose reach grows on a schedule that has been measured in months. Under those conditions the strategic variables are the rate at which a firm revises its theory, the quality of each revision, the share of revision that formalization cannot reach, and the coherence the firm can spend without ceasing to be one firm. This paper has put those variables in a model, shown that velocity has an optimum rather than a maximum, derived the half-life of derivable advantage from the formalization rate, treated the human edge as a declining empirical quantity with a floor whose value decides whether the theory describes durable advantage or its disappearance, and stated what observation would settle that question.

The model is a program, not a result. Its two exogenous parameters have proxies in the literature and its four endogenous ones have proposed estimators and no data. The next step is a panel of firms with a text corpus deep enough to estimate ω and Q, a classification protocol for λ, and a problem-class taxonomy for φ. Until that exists the paper's claims are claims about structure: that these are the right variables, that they interact in this way, and that the question of whether machines can theorize is a parameter to be measured rather than a premise to be defended.

Appendix A · Claims ledger

All quantitative and attributive claims in the paper, with the type and source of each. Sourced means traced to the primary document listed in the references during preparation of this edition. Model means computed from stated parameters in this paper. Interpretive means a reading of a philosophical text that the text does not itself make.

CLAIM	TYPE	SOURCE
Firms create value by forming theories that reveal opportunities invisible from data alone	SOURCED	Felin and Zenger (2017); Zenger (2013)
Present machine learning systems reason from data and cannot generate forward-looking causal beliefs beyond their training distribution	SOURCED	Felin and Holweg (2024); treated here as a hypothesis about λ_∞
Process philosophy has been applied to organizations and to strategy	SOURCED	Tsoukas and Chia (2002); Hernes (2014); Nonaka, Toyama and Hirata (2008); MacKay, Chia and Nair (2021)
Advantage as reconfiguration capacity; absorptive capacity as a rate	SOURCED	Teece, Pisano and Shuen (1997); Eisenhardt and Martin (2000); Cohen and Levinthal (1990)
Exploration and exploitation trade off; convex cost of revision	SOURCED	March (1991); functional form $\kappa\omega^2$ is this paper's assumption
The growth of knowledge cannot be predicted; the adjacent possible cannot be prestated	SOURCED	Popper (1957); Kauffman (2000); Felin, Kauffman, Koppl and Longo (2014); Knight (1921)
Citation half-lives: metallurgical engineering 3.9, physics 4.6, physiology 7.2, chemistry 8.1, geology 11.8 years	SOURCED	Burton and Kebler (1960), values as quoted in secondary summaries of the paper; the original article was not accessed in full during preparation

CLAIM	TYPE	SOURCE
Professional psychology half-life nearly 9 years, projected just over 7; range 3.3 (psychopharmacology) to 19 (psychoanalytic psychology)	SOURCED	Neimeyer, Taylor and Rozensky (2012), abstract
474 cirrhosis and hepatitis conclusions; 60 percent still true in 2000; half-life of truth 45 years; no survival advantage for stronger methodology	SOURCED	Poynard et al. (2002), abstract
50 percent task horizon doubling every ~7 months (212 days; 95% CI 171–249 days) since 2019; GPT-2 ~2 s; Claude 3.7 Sonnet ~50 min; o3 110 min; 170 tasks	SOURCED	Kwa et al. (2025), versions 1 and 3
Breakthrough patents identifiable by textual dissimilarity to prior and similarity to subsequent work	SOURCED	Kelly, Papanikolaou, Seru and Taddy (2021)
Engineering-degree half-life 35 years in 1930, ~10 in 1960	WITHDRAWN	No located primary study; trade-press origin; removed
Empirical evidence from 500 firms in four strategic quadrants; industry velocity trajectories 2015–2025	WITHDRAWN	June 2025 draft Figures 3 and 4 were hand-placed illustrations; no dataset existed; removed
Optima $\omega^* = 2, 5, 8$ and peaks 0.2, 1.25, 3.2 for $\lambda_{eff} = 0.2, 0.5, 0.8$	MODEL	Equation (5), $Q = 1$, $\kappa = 0.05$
Regime A frontier 30 → 11.6 units; Regime B 30 → 40.7 units over ten years	MODEL	Equations (1)–(3), parameters in Figure 3 caption
Whitehead: the many become one and are increased by one; Category of the Ultimate as creativity, many, one	SOURCED	Whitehead (1978), p. 21

CLAIM	TYPE	SOURCE
Bergson: intuition as intellectual sympathy placing oneself within an object; the universe as a machine for the making of gods	SOURCED	Bergson (1912); Bergson (1935), final sentence
K, D, N as objectified past, conformal feeling, conceptual feeling; metamorphic structure as an organizational analogue of the Category of Reversion	INTERPRETIVE	This paper's reading of Whitehead (1978); marked as extension
Execution runs in spatialized time; revision runs in duration	INTERPRETIVE	This paper's reading of Bergson (1910; 1911)

B ERRATA AGAINST THE JUNE 2025 DRAFT

Appendix B · Errata against the June 2025 draft

The draft's headline equation, Strategy = Epistemological Velocity × Human Edge, is replaced by equation (5). The draft's ω(t) = dΨ/dt × Q(t) had an undefined state function; Ψ is now defined as the firm's theory of value with a stated representation.

The draft's unit, paradigm shifts per year, is replaced by framework revisions per year.

The draft's premise that machine systems are confined to intellect and conformal feeling is withdrawn as a premise and restated as the hypothesis $\lambda_\infty > 0$.

The draft's Figure 1 caption dated the hepatology study to 2019 and the psychology poll to 2016; the studies are Poynard et al. (2002) and Neimeyer, Taylor and Rozensky (2012). Its engineering series had an IEEE attribution with no primary study and is removed. Its unsourced points for mathematics, physics at 15 years, chemistry at 25, software and machine learning are removed. Its required-velocity curve y = 10/x was arbitrary and is removed; the text called it a diagonal line when it was a hyperbola.

The draft's Figure 2 normalized the three regions to a fixed total and so could not show frontier expansion; its insight text used a different formula from its bars, and its static text stated a value of 2.0 where the formula gave 1.1. It is replaced by Figure 3 of this edition, which does not normalize.

The draft's Figure 3 was captioned as empirical evidence from 500 firms and consisted of twenty hand-placed points. Its Figure 4 presented invented industry trajectories with a caption citing analysis that did not exist. Both are removed.

The draft attributed to Bergson the sentence that the universe is a machine for creating novelty and a definition of intuition as the mind's direct participation in the universal becoming. Neither has been located. The attested sentences are given in Section 5.

The draft's references dated Creative Evolution to 1907 with the 1911 translation's publisher, Time and Free Will to 1889 with a 1910 English publisher, and Bergsonism to 1966 with its 1988 translator's publisher. The original date and the translation are now given separately for each. The draft cited Whitehead's Process and Reality from the 1929 printing; the corrected edition is used here. The draft's use of conceptual reversion is marked as an extension.

The draft claimed that strategic theory had privileged being over becoming and omitted the theory-based view, Felin and Holweg, process organization studies, dynamic capabilities, absorptive capacity, Knight, Popper and Kauffman. Section 2 supplies them.

The draft's earlier name for λ is replaced by human edge throughout this edition. The draft's quadrant label for high-ω, high-λ firms is not reused, since the quadrant chart it labeled was not data.

R REFERENCES

References

The entries were checked against publisher, index or repository records during preparation of this edition. Where the check reached only an index record rather than the document itself, the entry says so.

01 Arbesman, S. (2012). The Half-Life of Facts: Why Everything We Know Has an Expiration Date. New York: Current.

02 Argote, L. (2013). Organizational Learning: Creating, Retaining and Transferring Knowledge, 2nd ed. New York: Springer.

03 Bergson, H. (1910). Time and Free Will: An Essay on the Immediate Data of Consciousness, trans. F. L. Pogson. London: George Allen; New York: Macmillan. Original: Essai sur les données immédiates de la conscience, Paris: Alcan, 1889.

04 Bergson, H. (1911). Creative Evolution, trans. A. Mitchell. New York: Henry Holt. Original: L'Évolution créatrice, Paris: Alcan, 1907.

05 Bergson, H. (1912). An Introduction to Metaphysics, trans. T. E. Hulme. New York: G. P. Putnam's Sons. Original 1903.

06 Bergson, H. (1935). The Two Sources of Morality and Religion, trans. R. A. Audra and C. Brereton. New York: Henry Holt; London: Macmillan. Original: Les Deux Sources de la morale et de la religion, Paris: Alcan, 1932.

07 Burton, R. E., and Kebler, R. W. (1960). The “half-life” of some scientific and technical literatures. American Documentation, 11(1), 18–22. Index record checked; field values taken from secondary quotation.

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

09 Deleuze, G. (1988). Bergsonism, trans. H. Tomlinson and B. Habberjam. New York: Zone Books. Original: Le Bergsonisme, Paris: PUF, 1966.

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

11 Felin, T., and Holweg, M. (2024). Theory is all you need: AI, human cognition, and causal reasoning. Strategy Science, 9(4), 346–371.

12 Felin, T., Kauffman, S., Koppl, R., and Longo, G. (2014). Economic opportunity and evolution: Beyond landscapes and bounded rationality. Strategic Entrepreneurship Journal, 8(4), 269–282.

13 Felin, T., and Zenger, T. R. (2017). The theory-based view: Economic actors as theorists. Strategy Science, 2(4), 258–271.

14 Hernes, T. (2014). A Process Theory of Organization. Oxford: Oxford University Press.

15 Kauffman, S. A. (2000). Investigations. New York: Oxford University Press.

16 Kelly, B., Papanikolaou, D., Seru, A., and Taddy, M. (2021). Measuring technological innovation over the long run. American Economic Review: Insights, 3(3), 303-320.

17 Knight, F. H. (1921). Risk, Uncertainty and Profit. Boston: Houghton Mifflin.

18 Kuhn, T. S. (1962). The Structure of Scientific Revolutions. Chicago: University of Chicago Press.

19 Kwa, T., West, B., Becker, J., Deng, A., Garcia, K., Hasin, M., Jawhar, S., Kinniment, M., Rush, N., Von Arx, S., Bloom, R., Broadley, T., Du, H., Goodrich, B., Jurkovic, N., Miles, L. H., Nix, S., Lin, T., Painter, C., Parikh, N., Rein, D., Sato, L. J. K., Wijk, H., Ziegler, D. M., Barnes, E., and Chan, L. (2025). Measuring AI ability to complete long software tasks. arXiv:2503.14499. Version 1 (March 2025) and version 3 consulted; the title of version 1 omits “software”.

20 MacKay, B., Chia, R., and Nair, A. K. (2021). Strategy-in-practices: A process philosophical approach to understanding strategy emergence and organizational outcomes. Human Relations, 74(9), 1337-1369.

21 March, J. G. (1991). Exploration and exploitation in organizational learning. Organization Science, 2(1), 71-87.

22 Neimeyer, G. J., Taylor, J. M., and Rozensky, R. H. (2012). The diminishing durability of knowledge in professional psychology: A Delphi poll of specialties and proficiencies. Professional Psychology: Research and Practice, 43(4), 364-371.

23 Nonaka, I., Toyama, R., and Hirata, T. (2008). Managing Flow: A Process Theory of the Knowledge-Based Firm. Basingstoke: Palgrave Macmillan.

24 O'Reilly, C. A., and Tushman, M. L. (2004). The ambidextrous organization. Harvard Business Review, 82(4), 74-81.

25 Popper, K. R. (1957). The Poverty of Historicism. London: Routledge & Kegan Paul.

26 Poynard, T., Munteanu, M., Ratziu, V., Benhamou, Y., Di Martino, V., Taieb, J., and Opolon, P. (2002). Truth survival in clinical research: An evidence-based requiem? Annals of Internal Medicine, 136(12), 888-895.

27 Rescher, N. (1996). Process Metaphysics: An Introduction to Process Philosophy. Albany: State University of New York Press.

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

29 Tsoukas, H., and Chia, R. (2002). On organizational becoming: Rethinking organizational change. Organization Science, 13(5), 567-582.

30 Whitehead, A. N. (1938). Modes of Thought. New York: Macmillan.

31 Whitehead, A. N. (1978). Process and Reality: An Essay in Cosmology, corrected edition, ed. D. R. Griffin and D. W. Sherburne. New York: Free Press. Original 1929.

32 Zenger, T. R. (2013). What is the theory of your firm? Harvard Business Review, 91(6), June 2013, reprint R1306D.

EPISTEMOLOGICAL VELOCITY · REVISED WORKING PAPER · AUGUST 2026 · BRAD HEITMANN

Set in Cormorant Garamond, Inter and JetBrains Mono. Figures are inline SVG generated from stated data and parameters; the generator and its sources are listed in Appendix A.