Distributed Genius and the Generative Ratchet: How Inherited Possibility Becomes Individual Capability, and When It Only Appears To

PIATRA . INSTITUTE

July 2026

Abstract

Genius is usually located in a person. This paper locates most of its cultural work in the environment a learner inherits, the exemplars, tools, and possibility-spaces that exceptional achievements leave behind, and treats education as the process that converts that distributed inheritance into individual capability. The move reframes the pedagogical problem. Cumulative-culture theory explains how a civilization keeps its gains across generations through a ratchet; it does not explain how a kept gain becomes an individual's usable and extendable skill, and that conversion is what a theory of education has to describe. The paper's operational core is a distinction among three frontiers. The production frontier is the best output a learner can put on the page with the tools at hand; the competence frontier is what the learner can independently generate and revise; the judgment frontier is what the learner can reliably evaluate and select. Generative AI expands production almost immediately while leaving competence and judgment near their starting values, and the resulting gap is false ratcheting, sophisticated output unbacked by durable capability, revealed the moment the tool is withdrawn. The distinction is measurable, which is what separates the framework from a metaphor: test the learner with the tool present and again with it absent, and read the difference. An accompanying illustrative model makes the logic explicit. Under one learning dynamics, a finished-output tool and a structured tool that fades both reach the same production frontier near 0.87, yet competence ends at 0.15 in the first case and 0.84 in the second; hand every learner the tool at a common strength and then remove it, and displayed output collapses by 83.1% for the finished-output tool against 10.3% for the fading one; the genuine-ratchet fraction, the share of the production advance that became durable capability, is 0.10 against 0.99, and the ordering survives 2,000 seeded parameter perturbations. The model is illustrative, its parameters stipulated to instantiate the definitions rather than measured from learners. The framework then states its predictions as falsifiable claims: AI ratchets capability genuinely when it decomposes an achievement, elicits prediction and explanation before revealing answers, generates controlled contrasts, preserves productive difficulty, and withdraws its own assistance, and it counterfeits capability when it mainly supplies finished products. A single randomized trial of a heavily scaffolded physics tutor reports a large short-term gain of 0.73 to 1.3 standard deviations (Kestin et al., 2025), and a field experiment finds an unrestricted chatbot raising practice performance while lowering later unaided scores (Bastani et al., 2025). The framework predicts that split and specifies the measurements that would confirm or break it.

Genius Was Never Only in the Genius

A contemporary student of composition can write counterpoint that would have taken Bach's contemporaries years to reach, and can do it in a first term. Nothing about the student's raw endowment explains this. What explains it is everything that arrived between: the tempered keyboard, figured-bass pedagogy, the printed chorales, two centuries of theory that named what Bach did so it could be taught. The capability that once lived inside one exceptional person now lives, in part, in the apparatus that surrounds an ordinary one. The student inherits the results of a search without paying its cost.

Read this way, the exceptional individual is one component of a larger productive system rather than its sole seat. Csikszentmihalyi (1999) argues that creativity is a property of a system with three interacting parts, the individual who varies, the domain that stores symbolic rules, and the field that selects what enters the domain, so that a contribution is creative only when the field admits it and the domain carries it forward. Cognition itself has the same shape. Hutchins (1995) shows a navigational fix computed not in any one sailor's head but across instruments, charts, procedures, and a division of labour, a cognitive process in its own right whose parts are people and artifacts. Clark and Chalmers (1998) push the point to its edge with active externalism: when an external resource is coupled tightly enough to the reasoning that uses it, it functions as part of the mind rather than as a passive input to it. Capability, on all three accounts, is distributed across a standing arrangement of people and things.

The individual-effort tradition seems to cut the other way, and its own evidence does not. Ericsson, Krampe and Tesch-Römer (1993) tie expert performance to deliberate practice, the effortful, feedback-rich, individually targeted training that separates experts from the merely experienced. Deliberate practice is real and it is indispensable. It is also insufficient as an account of who reaches the frontier. Macnamara, Hambrick and Oswald (2014), meta-analysing the practice literature, find that accumulated practice explains a modest and sharply domain-dependent share of performance variance, on the order of a fifth to a quarter in music and games and far less in education and the professions. The residual is not mystery. Much of it is the environment: which exemplars a learner meets, how early, decomposed by whom, with what tools for varying them. Henrich (2016) makes the population-level version of the claim, that human success rests less on individual brainpower than on a collective brain, the accumulated and transmitted know-how that no single lifetime could reconstruct.

Call the resulting object distributed genius: the spread of exceptional generative capability across artifacts, notations, practices, institutions, tools, and learners, such that achievements once dependent on rare individuals become navigable and extendable by many. The term is meant literally and narrowly. It does not claim that groups are smarter than people, nor that individuals cease to matter. It claims that the generative power a culture attributes to its geniuses is substantially carried by the infrastructure those geniuses left behind, and that this infrastructure is what education hands to a learner. The question a theory of education must then answer is not how rare individuals produce exceptional work. It is how the standing infrastructure becomes a particular person's capability.

How an Exemplar Restructures the Possible

The word inheritance undersells what a transformative work does. A masterpiece is not only a point in a domain that later people admire. It changes the shape of the domain, so that positions unreachable before it become reachable after.

Write Ω_t for the possibility-space of a domain at time t, the configurations that its practitioners can currently conceive, attempt, and recognize as belonging to the domain. A merely excellent work occupies an open point in Ω_t. A transformative exemplar E does something stronger. It enlarges the space:

$$\Omega_{t+1} \;=\; \Omega_t \,\cup\, D(E) \,\cup\, V(E) \,\cup\, Q(E) \,\cup\, R(E, \Omega_t),$$

where $D(E)$ is the set of operations the work demonstrates can be done, $V(E)$ the variations it makes newly visible, $Q(E)$ the questions it raises that did not exist before it, and $R(E, \Omega_t)$ the recombinations it enables with the rest of the culture already present. The union is the content of the claim that a work opens a space rather than filling a slot in one.

Each term names a distinct kind of work the exemplar performs. Demonstration converts an uncertainty into a settled possibility: after Michelangelo, the monumental torqued body is known to be drawable, and a later artist begins from that certainty rather than from doubt. Search compression banks the cost of the failures the exemplar already ran, so that successors inherit the outcome of an exploration without repeating it. Perceptual reorganization retrains what practitioners can see; a torso, after certain drawings, shows tensions and masses that were physically present all along and perceptually unavailable. New vocabulary supplies forms that can be named, varied, and recombined, the fugue as a manipulable object rather than a single piece. Problem creation opens questions the work leaves standing, which become the frontier the next generation works. Euler's manipulations of infinite series demonstrated operations, compressed a century of trial, and left the questions about convergence that the next century had to answer. The adjacent possible, in Kauffman's (2000) sense, is just this expanding rim of states reachable in one step from where the culture now stands, and a transformative work is a large, deliberate expansion of it.

None of this makes the enlargement automatically available to any given learner. The space Ω_{t+1} is public. A person standing in front of the drawing sees a finished figure, not the decision sequence that produced it; the operations in $D(E)$ are demonstrated but not thereby transferred. Vygotsky (1978) named the general form of the gap and its closing: higher capabilities exist first between people, in externalized cultural tools and signs, and become individual only through a process of internalization that social interaction has to drive. The masterpiece externalizes a structure. Whether that structure becomes anyone's capability is a separate question, and it is the pedagogical one.

The Ratchet Preserves the Gain but Withholds the Capability

Cumulative culture has a canonical mechanism for keeping enlargements once they occur. Tennie, Call and Tomasello (2009) call it the ratchet: human populations retain modifications and build on them across generations, where other species' traditions stay within a zone of solutions each individual could reinvent. The retention depends on high-fidelity transmission. Tomasello, Kruger and Ratner (1993) locate that fidelity in distinctively human forms of cultural learning, imitative, instructed, and collaborative, that copy not just outcomes but the intent and procedure behind them. Boyd and Richerson (1985) supply the population dynamics under which such transmitted variation accumulates into adaptations no individual designed.

The ratchet is a theory of preservation, and preservation is a different problem from acquisition. That a technique is retained in the collective record guarantees that the next generation can find it. It says nothing about whether any particular member of that generation can wield it. A library holds the calculus; a student does not thereby know it. The ratchet keeps Ω_{t+1} from decaying back toward Ω_t. It leaves open the question of how a point in the preserved space becomes an operation in a specific learner's hands.

This is the seam the present framework works. Between the collective and the individual sits a conversion that the cumulative-culture literature brackets, because its unit of analysis is the population and its currency is the trait's frequency across it. Education is the institution built for that conversion, and its success condition is not that the culture retains an achievement but that a learner ends able to generate, vary, and extend it without the scaffolding that taught it. Henrich's (2016) collective brain explains why the species is capable far beyond any of its members. It does not, and does not try to, explain how to make a given member capable. The conversion has its own dynamics, its own failure modes, and, as the next sections argue, its own measurements, and generative AI acts directly on it.

Three Frontiers: Production, Competence, Judgment

Assessment usually reports one number, the quality of what a learner produced. When the learner worked with a capable tool, that number conflates capacities that a theory of the conversion must hold apart.

Definition 1. *For a learner i on a domain of tasks, the* production frontier P_i *is the highest quality of output i can produce using all tools currently available. The* competence frontier C_i *is what i can independently generate, reconstruct, diagnose, and revise without those tools. The* judgment frontier J_i *is what i can reliably evaluate: to tell a stronger output from a weaker one, detect an error, recognize cliché, and defend a selection.*

The three come apart because they are supported differently. Production can be borrowed; a learner wielding a strong model produces at the model's level, not their own. Competence cannot be borrowed, by construction, because it is what remains when the borrowing stops. Judgment is partly independent of both: a person can sometimes tell that an output is wrong without being able to produce a right one, and, more dangerously, can produce fluent output while unable to tell whether it is any good. Competence and judgment are coupled, since reliable evaluation of generative work generally requires enough internal structure to know how the work is made, but the coupling is loose enough that the two frontiers move at different rates and have to be tracked separately.

The distinction is not extended cognition renamed. Clark and Chalmers (1998) are right that a coupled tool can be part of the cognitive system that produces an output; the production frontier grants exactly that. The educational question is orthogonal to the metaphysical one. Whatever we say about where the mind stops, a learner who can perform only while coupled to a particular tool has acquired something different from a learner who can perform when the coupling is cut, and the difference is what schooling is for. Competence is the name for what survives decoupling. A pedagogy indifferent to it optimizes the wrong frontier.

Generative AI drives a wedge between the frontiers because it lifts production immediately and competence slowly, and it lifts production most for those with the least competence. Brynjolfsson, Li and Raymond (2025), studying customer-support agents given a generative assistant, find an average productivity gain concentrated among the least experienced workers, whose output moves closest to that of the best. Read as a fact about production, this is the technology working as intended. Read as a fact about the frontiers, it is a warning: the tool most enlarges the distance between what a novice can produce and what a novice can do. The central prediction of the framework follows. Where the interaction supplies output rather than building capacity, the change in production over a course of study greatly exceeds the change in competence and judgment,

$$\Delta P_i \gg \Delta C_i,\ \Delta J_i,$$

and the gap is invisible in any assessment that reads production alone. It becomes visible under one manipulation, removing the tool, which is what makes the framework testable rather than evocative.

Generative Regeneration: The Fixed Exemplar Becomes a Field

Everything so far predates the technology. Exemplars have always opened spaces, education has always faced the conversion, the frontiers have always been distinct in principle. Generative models change one thing, and it is consequential: they make the exemplar interactive.

A printed chorale is fixed. A learner can copy it, annotate it, and study it, but it does not answer. A generative model standing next to the same chorale can re-voice it in another style, strip out a suspension to show what the suspension was doing, transpose the problem to an instrument that did not exist when it was written, generate the piece the composer might have written next, and produce a plausibly defective version for the learner to diagnose. Call this generative regeneration: the reconstruction of a fixed cultural achievement as a variable, interactive, and counterfactual field rather than a static object. The exemplar stops being a point to be admired and becomes a neighbourhood to be explored, its decisions individually adjustable, its invariants exposed by watching what breaks when a decision is changed.

This is the active form of the possibility-space operator. The terms $D(E)$, $V(E)$, and $R(E, \Omega_t)$ were, for most of history, latent: the variations an exemplar made visible still had to be executed by hand, one at a time, at the cost of the skill required to execute them. Regeneration makes them queryable on demand and at near-zero marginal cost. The neighbourhood of an achievement, the family of nearby works it implied but never contained, becomes locally explorable by someone who could not yet produce a single member of it unaided. Where the classical ratchet preserved one solution and handed it forward, a generative ratchet can hand forward the surrounding field of solutions, and can do so fast.

Speed is the promise and the trap. The same mechanism that lets a learner traverse in an afternoon a space that once took an apprenticeship can also traverse it for the learner, depositing the destination without the journey. Regeneration decomposes an achievement into its decisions, and it can also skip the decomposition and return the finished object. It can generate a controlled contrast that forces an explanation, and it can generate the explanation too, so that nothing is left for the learner to do. Cognitive apprenticeship, in the sense of Collins, Brown and Newman (1989), works by making an expert's normally hidden process visible so a learner can take it over, through modelling, coaching, and the gradual handover of responsibility. Generative regeneration can serve that handover or foreclose it, and which it does is a property of how the interaction is arranged rather than of the model. The next section makes the fork quantitative.

False Ratcheting, Measured

To show that the three-frontier distinction is coherent and separates the cases it names, consider a deliberately simple model of a course of practice. The model is illustrative. Its parameters are stipulated to instantiate the definitions above rather than estimated from learners, and its numbers are facts about the model's logic rather than measurements of teaching. What it establishes is that the framework's predictions follow from its definitions under a transparent dynamics, and that the tool-removal test discriminates the regimes the framework says it should.

A learner carries a competence frontier C and a judgment frontier J, each running from 0.10 at intake toward a ceiling of 1. Production is an algebraic function of competence and the assistance a currently supplied, $P = C + a\,(\kappa - C)$, with tool ceiling $\kappa = 0.95$, so that at full assistance the learner produces near the tool's level whatever their own, and at zero assistance production falls back exactly to competence. Over thirty practice sessions a single update law moves C and J. Competence grows only to the extent that the learner does the generative work themselves, receives feedback that corrects it, and operates at a productive level of difficulty, a managed gap above current competence that is neither trivial nor overwhelming; a desirable-difficulty term, peaked at that gap and vanishing when the task is trivial, encodes the effect that Bjork and Bjork (2011), Kapur (2008), and Sweller (1988) document from three directions, that learning is largest at an intermediate, effortful challenge and smallest when the work is either done for the learner or left crushingly unsupported. Judgment grows with the evaluative work the interaction demands, weighted by the competence needed to make it reliable.

Three regimes supply the same dynamics with different interaction signals. The no-tool baseline gives the learner all of the work but leaves them facing the raw, unscaffolded gap. The finished-output regime holds assistance high and constant, has the tool do most of the work and supply the answer rather than the operations, and asks for little evaluation, the arrangement of an unrestricted chatbot used to complete a task. The structured-fading regime starts with high assistance and withdraws it to zero across the course, has the tool structure the work rather than perform it, demands prediction and evaluation at each step, and holds the learner at the productive gap, the arrangement of a scaffolded tutor that hands responsibility back.

The regimes diverge as the framework predicts. Both tools reach essentially the same production

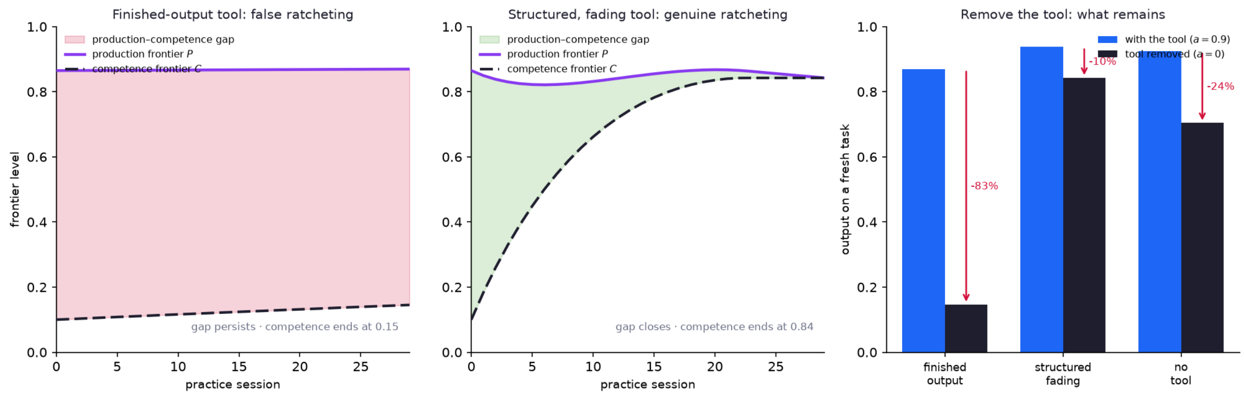

Figure 1: Three frontiers over a course of practice. Left, a finished-output tool holds production near the ceiling while competence barely moves, and the gap between them (shaded) persists to the end: false ratcheting. Centre, a structured tool that fades lets competence rise to meet production, and the gap closes as assistance withdraws: genuine ratcheting. Right, the tool-removal test, giving every regime the tool at a common strength and then taking it away, reads how much displayed output was really the learner's; the finished-output learner loses most of it. The numbers are properties of an illustrative model rather than measurements.

frontier, 0.87, so any assessment reading output alone would rate them equivalent. Their competence frontiers end far apart: 0.15 under the finished-output tool, 0.84 under the structured tool that fades, and 0.70 for the learner who used no tool at all. Judgment separates the same way, ending at 0.23 against 0.89. The middle comparison is the sharp one. The finished-output learner ends less competent than the learner who never had a tool, while looking more accomplished than either all term, so a tool that maximizes visible output can leave a learner below the level of unaided practice.

The gap is latent until the tool is removed. Give each learner the tool at a common strength and then withdraw it, and production falls to whatever competence backs it. For the finished-output learner, displayed output collapses by 83.1%; for the structured-fading learner, by 10.3%; for the no-tool learner, by 23.9%, all of it the honest cost of losing a tool they had learned to work with rather than through. The finished-output regime collapses about 8 times as much as the structured one. Summed into the incorporation ratio the framework proposes, the share of the production advance that became durable capability, weighting judgment at $\lambda = 0.5$, the genuine-ratchet fraction is 0.10 for the finished-output tool and 0.99 for the structured one. The first is false ratcheting made numerical: production advanced almost a full unit and almost none of it stayed with the learner. The second is genuine ratcheting, where the production the learner ends up showing is production they can now generate themselves.

The ordering is not an artifact of one parameter set. Perturbing the model's constants by 15% in either direction, across 2,000 draws under a fixed seed, the finished-output regime ends below the structured one on competence in every draw, and collapses more on removal in every draw. The separation is a feature of the interaction architecture rather than a knife-edge choice of numbers. What the model cannot do is say how large the real gap is in a real classroom; its parameters were chosen to express the definitions, and only measurement can fill them. It shows that the distinction is coherent, that it has a determinate empirical signature, and that the signature is the behaviour on tool removal.

What Makes a Tool Ratchet Genuinely

The model localizes the difference between the regimes in their interaction signals, which is where design lives. Turning the contrast into guidance yields a set of conditions, each stated as a claim that a study could break.

A tool ratchets genuinely when it decomposes an achievement into the decisions that produced it rather than returning the achievement whole. The prediction is an interaction, not a main effect: decomposition should raise unaided transfer more than it raises assisted output, so a design that improves the finished product least may improve durable competence most. Cognitive apprenticeship rests on this externalization of hidden process (Collins, Brown, and Newman, 1989), and the framework locates its benefit on the competence frontier, not the production one.

It ratchets, too, when it draws a prediction and an explanation out of the learner before it reveals an answer. Requiring a commitment to what will happen, and why, before the model shows what does, turns a passive display into a test the learner takes. The principle is that AI helps when it redirects effort toward comparison, diagnosis, and explanation, and harms when it removes the operation the learner was meant to acquire.

Controlled contrast is a third lever. A single finished example demonstrates possibility; a family of variations that differ in one decision at a time teaches what each decision does, but only if the learner predicts and explains the differences instead of scrolling past them. The tool generates the contrasts and leaves their reading to the learner.

Difficulty is the fourth. The desirable-difficulties (Bjork and Bjork, 2011) and productive-failure (Kapur, 2008) literatures agree, from opposite starting points, that struggle at an appropriate level builds durable understanding, and cognitive-load theory (Sweller, 1988) marks the other boundary, where difficulty without support becomes overload. A tool that smooths every difficulty to zero optimizes the learner out of the learning.

The last condition is withdrawal. Scaffolding, in the original sense of Wood, Bruner and Ross (1976), is defined by its removal: support is given for what the learner cannot yet do and taken away as they become able, and a scaffold that never comes down is a permanent prosthesis. The design goal is not that the learner never uses the tool. It is that the learner can function when the tool is gone.

The field evidence, still thin, already has the shape of the prediction. Bastani et al. (2025), in a field experiment with roughly 1,000 secondary-school mathematics students, found that access to an unrestricted model raised performance by about half during assisted practice and then left students scoring below their tool-free peers once the model was taken away, while a guardrailed version that guided rather than answered removed the harm. The two arms are the model's two tool regimes in a real classroom, and they came apart on removal as the framework says they must. On the other side, Kestin et al. (2025) report a randomized trial in which a carefully engineered physics tutor, built with scaffolding, step-by-step guidance, and targeted feedback, produced a large learning gain estimated at 0.73 to 1.3 standard deviations over in-class active learning, in less time. The result is a single suggestive data point and the authors say so: it measures immediate post-lesson outcomes rather than durable retention, uses a heavily prompt-engineered tutor rather than a generic chatbot, and covers two physics topics in one course. It does not validate the framework. It is consistent with the framework's sharper claim, which is not that AI tutoring works but that its interaction architecture is the variable deciding whether it does, so that a scaffolded tutor and an answer-dispensing chatbot are two different treatments with opposite signs rather than two doses of one.

Two collective effects extend the same logic past the individual learner. Doshi and Hauser (2024) find that AI-generated ideas raise the assessed creativity of individual stories, most for the least creative writers, while making the stories more similar to one another, so that optimizing each learner's output can narrow the cohort's diversity even as it lifts the average. Recursive dependence on generated material carries a distribution-level risk of its own: Shumailov et al. (2024) show that models trained on their own output collapse toward the head of the distribution and lose its tails. Their result is about models, and no equivalent process in human culture has been demonstrated, but it supplies a precise analogy and a testable worry. A pedagogy that lets generated approximations replace the primary human archive, rather than anchor exploration to it, risks teaching the mean of the model in place of the range of the tradition.

The Measurement That Decides It

The framework matters only if it changes what is measured, and it prescribes one change above the rest: stop treating assisted output as the criterion of learning. The decisive quantities are what the learner can do when the tool is switched off, and they are directly observable.

construct	operational measure
assisted production	quality produced with the tool available
autonomous competence	quality produced after the tool is removed
assistance gap	assisted minus unassisted, on matched tasks
durable competence	unassisted performance after a delay of weeks
transfer	unassisted performance on a structurally new task
judgment calibration	agreement of the learner's selections and confidence with expert assessment
error detection	rate of catching planted or model-generated defects
collective diversity	dispersion of outputs across a cohort

The assistance gap is the practical form of the tool-removal test, and it turns the paper's central claim into a study anyone can run. Assign learners to a finished-output tool, a structured tool that fades, and no tool; equate the tasks; and measure not what each group produces with help but what each can produce and evaluate without it, immediately and after a delay. The framework predicts that assisted output will fail to order the groups and that unassisted competence and judgment will separate them, with the fading tool ahead, the finished-output tool behind, and the gap widening with delay. A single trajectory, a learner working an exemplar into an artifact and then into a transferable capability, holds the evidence that any finished artifact hides, which is why assessment has to follow the path and not only inspect the endpoint.

The same design says how the framework fails. If adequately powered trials find that unassisted competence and judgment track assisted production regardless of how the interaction is arranged, so that finished-output and structured-fading tools leave learners equally capable once the tool is gone, the three-frontier distinction collapses and the framework is wrong. If fading and scaffolding show no advantage in delayed unaided transfer over tools that simply supply the answer, the design conditions of the previous section are empty. These are not hedges. They are the observations that would retire the account, and stating them is the price of claiming that it is more than a redescription of what good teachers already know.

That redescription charge is worth meeting directly, because much of the machinery here is old. Cumulative culture supplies the ratchet, extended cognition the distributed ontology, cognitive apprenticeship and desirable difficulties and productive failure the design conditions. What is new is small and, if it holds, decisive: the separation of production from competence and judgment into frontiers that move at different rates under a tool, and the recognition that their divergence is measurable on removal. The neighbouring theories were built for a world in which producing an output at a given level was itself evidence of the capability to produce it, because production was expensive and could not be borrowed. Generative AI breaks that link. When output becomes nearly free and fully borrowable, the quantity that used to certify learning stops certifying anything, and the scarce, decisive act in education inverts. For most of its history teaching worked by supplying what the learner lacked. When supply is what the tool does best and for almost nothing, the teaching move that stays scarce is the withdrawal of assistance, the deliberate return of difficulty, the refusal to act for the learner at the points where acting for them would feel most helpful. Distributed genius becomes an individual's capability only through an environment that sometimes declines to be genius on their behalf. The measure of a pedagogy, in the generative age, is what is left standing in the learner when the tool is turned off.

References

Bastani, H., Bastani, O., Sungu, A., Ge, H., Kabakcı, Ö., & Mariman, R. (2025). Generative AI without guardrails can harm learning: Evidence from high school mathematics. *Proceedings of the National Academy of Sciences*, 122(26), Article e2422633122.

Bjork, E. L., & Bjork, R. A. (2011). Making things hard on yourself, but in a good way: Creating desirable difficulties to enhance learning. In M. A. Gernsbacher, R. W. Pew, L. M. Hough, & J. R.

Pomerantz (Eds.), *Psychology and the real world: Essays illustrating fundamental contributions to society* (pp. 56–64). Worth Publishers.

Boyd, R., & Richerson, P. J. (1985). *Culture and the evolutionary process*. University of Chicago Press.

Brynjolfsson, E., Li, D., & Raymond, L. R. (2025). Generative AI at work. *The Quarterly Journal of Economics*, 140(2), 889–942.

Clark, A., & Chalmers, D. (1998). The extended mind. *Analysis*, 58(1), 7–19.

Collins, A., Brown, J. S., & Newman, S. E. (1989). Cognitive apprenticeship: Teaching the crafts of reading, writing, and mathematics. In L. B. Resnick (Ed.), *Knowing, learning, and instruction: Essays in honor of Robert Glaser* (pp. 453–494). Lawrence Erlbaum Associates.

Csikszentmihalyi, M. (1999). Implications of a systems perspective for the study of creativity. In R. J. Sternberg (Ed.), *Handbook of creativity* (pp. 313–335). Cambridge University Press.

Doshi, A. R., & Hauser, O. P. (2024). Generative AI enhances individual creativity but reduces the collective diversity of novel content. *Science Advances*, 10(28), Article eadn5290.

Ericsson, K. A., Krampe, R. T., & Tesch-Römer, C. (1993). The role of deliberate practice in the acquisition of expert performance. *Psychological Review*, 100(3), 363–406.

Henrich, J. (2016). *The secret of our success: How culture is driving human evolution, domesticating our species, and making us smarter*. Princeton University Press.

Hutchins, E. (1995). *Cognition in the wild*. MIT Press.

Kapur, M. (2008). Productive failure. *Cognition and Instruction*, 26(3), 379–424.

Kauffman, S. A. (2000). *Investigations*. Oxford University Press.

Kestin, G., Miller, K., Klales, A., Milbourne, T., & Ponti, G. (2025). AI tutoring outperforms in-class active learning: An RCT introducing a novel research-based design in an authentic educational setting. *Scientific Reports*, 15, Article 17458.

Macnamara, B. N., Hambrick, D. Z., & Oswald, F. L. (2014). Deliberate practice and performance in music, games, sports, education, and professions: A meta-analysis. *Psychological Science*, 25(8), 1608–1618.

Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024). AI models collapse when trained on recursively generated data. *Nature*, 631, 755–759.

Sweller, J. (1988). Cognitive load during problem solving: Effects on learning. *Cognitive Science*, 12(2), 257–285.

Tennie, C., Call, J., & Tomasello, M. (2009). Ratcheting up the ratchet: On the evolution of cumulative culture. *Philosophical Transactions of the Royal Society B*, 364(1528), 2405–2415.

Tomasello, M., Kruger, A. C., & Ratner, H. H. (1993). Cultural learning. *Behavioral and Brain Sciences*, 16(3), 495–511.

Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes* (M. Cole, V. John-Steiner, S. Scribner, & E. Souberman, Eds.). Harvard University Press. (Original works written 1930–1934.)

Wood, D., Bruner, J. S., & Ross, G. (1976). The role of tutoring in problem solving. *Journal of Child Psychology and Psychiatry*, 17(2), 89–100.