



The Best Teacher Is Now Free.


Learning Didn't Explode.





9% Usage Hit. Over a decade of free access, yet only 9% of students hit Khan Academy's own recommended usage targets (2024 Efficacy Report).


10,000+ students (Grades 3-6)
at 35 mins/week = +0.12 to 0.17
Standard Deviations in math.



7th/8th graders averaging
10 mins/week = 0.00 Gain.

You can give the machine away, but you cannot give away the desire to use it.

Scarce Books



System:
Master & Apprentice

Core Skill:
Memorization


Scarce Teachers



System:
Batch-Processing
(Prussian Model Classrooms,
Timetables, Age-slicing)

Core Skill:
Standardization

Abundant Instruction (Marginal cost = Zero)



System:
AI Tutors
(Always-on, infinitely patient)

New Scarce Resource:
Motivation

**The bottleneck didn't disappear. It moved from
'Who will teach?' to 'Why would anyone want to learn?'**

The Machine Burns The Very Resource We Now Need.

The 1973 Stanford Certificate Study


Preschoolers who love drawing + Promised a tangible certificate.

Result: Drawing rates dropped by 50% during free-play. (Confirmed by meta-analysis of 128 subsequent experiments).


The mechanism: People infer motives from behavior. If I draw for the certificate, the certificate becomes the only reason to draw.

The bad trade: School batch-distributes instruction while grinding down motivation with external carrots and sticks.

The Myth: Finding Demand




The Reality: Igniting Demand




2009 Science RCT:

High schoolers wrote one paragraph weekly: "How does what you learned relate to your life?" Both grades and interest rose. Crucially, the largest gains were among low-motivation students expecting to fail.

Building The Arena: The Architecture of Ignition.





The Geometry of Reward Shaping


The Failure



Sparse Distal Reward

No signal at intermediate steps.
A classic dead end.

The Fix




Reward Shaping (1999 RL Theorem).

Lay a trail of immediately redeemable rewards. The optimal policy doesn't change, but the gradient is paved.

Example: "Learn AI programming" (Distal) vs.
"Build a web game your friends can play today" (Shaped).


The AI Trap: Practice Scores Up 48%, Exam Scores Down 17%

Data Source: 2025 PNAS Study (1,000 Turkish High Schoolers, Math)



Full motivation plus an answer-giving AI just makes a student burn burn faster. They outsource the labor and bypass learning entirely.

Smooth learning feels fast, but nothing accumulates.




The Cognitive Law of Friction: The operations that feel effortful are the very process that carves information into long-term memory.



Desirable Difficulties. Spacing, self-testing, and struggling before seeing the answer.



AI is the most powerful difficulty-remover humanity has ever built. Outsourcing cognitive labor feels like learning at superhuman speed, but leaves the foundation empty.



Every sign of efficacy is decided by one variable: whether the cognitive labor of generating the answer still lives in the student.

The Inequality Warning: You Cannot Want What You Cannot Imagine.


Need-driven learning is structurally a greedy search.

It only climbs the gradients it can currently perceive. A student who has never touched calculus will never *need* calculus.

Science 2015: 68 free Harvard/MIT courses skewed significantly wealthier than the US average. Free access handed dividends to those who already possessed drive, priors, and self-regulation.

Pull is the dividend of **push**. Motivation activates an existing capability; it cannot build the foundation from zero.

Society's Extended Time Horizon.



The Child's Account:

Denominated entirely in immediate preferences. Controlled by the hyperbolic discount curve.


The Education Check:

'You'll need it someday' is a check written against reflective preferences. It cannot be cashed on the spot.

The System's True Role:


Forced exposure and liberal education build the priors before the need arrives, so that future needs become representable at all. A pull-only education hands the curriculum over to a steep discount curve.

AI Alignment (Preventing Value Drift)




The engineering goal is to keep an agent stably aligned to an external objective. Drifting or rewriting a reward function (wireheading) is catastrophe. Domestication.

Education (Controlled Value Drift)



A student growing an interest she didn't have, rewriting what she wants. What counts as disaster in AI alignment is called graduation in education.

The Double Alignment Blueprint.



The goal is not an agent aligned to a machine. The goal is a **human aligned to their own capability.**