ESSAY 3

THE TIMING PARADOX

Why We Train People at Exactly the Wrong Moment

SHAURAV SEN

This essay is part of
Training Is Broken: Learning Doesn't Have to Fail.
Download the complete book and all 12 essays at
shaurav.org

ESSAY 3

The Timing Paradox

Why We Train People at Exactly the Wrong Moment

We design learning calendars around what's convenient for us, not what works for those who need to learn—and then blame learners when nothing sticks.

SYNOPSIS:

We design learning around calendars, not moments of need. We send people to workshops months before they'll use the skill, flood them with content during onboarding, and roll out annual programs tied to budget cycles. Then we act surprised when nothing sticks. Cognitive science has told us for over a century that timing is everything: people retain more when learning is close to application and reinforced over time. The paradox? Organizations cling to scheduled training because it's administratively convenient, while employees quietly solve real problems through just-in-time, self-directed learning. In the age of AI, this gap becomes impossible to ignore. The question isn't "Can we time learning perfectly?" It's: *How much closer can we get—and what would it look like if timing was driven by work, not calendars?*

When the Moment Finally Arrives

In January, a new manager attends a three-day leadership program. She practices coaching conversations, learns a feedback model, and leaves with a thick workbook and a 4.8 satisfaction score. In July, she faces her first truly hard conversation. An employee is missing deadlines, the team is frustrated, and her own boss is asking pointed questions.

She freezes.

The models from January are gone. What she has instead is a racing heart and a browser tab: *"how to give tough feedback to defensive employee."*

She didn't need the training when it was delivered. She needs it *now*—when the stakes are real, emotions are high, and her brain is fully awake. That is the timing paradox in one scene: we schedule learning when it suits the organization; people need it when it collides with reality.

Just-in-Case vs Just-in-Time

Most corporate training runs on a just-in-case logic: teach this now, **just in case** it's needed one day. Front-load everything into onboarding, **just in case** it matters later. Run the workshop this quarter, **just in case** next year's budget is cut.

It shows up everywhere—annual compliance marathons that bear little resemblance to real ethical dilemmas, new-hire "bootcamps" that compress months of knowledge into a week, leadership programs scheduled around vendor calendars instead of role transitions.

On paper, this looks responsible: prepare people in advance. But human memory doesn't work that way.

The forgetting curve, first described by Hermann Ebbinghaus, shows that without reinforcement or use, we lose most new information in days

or weeks. When you teach people something months before they need it, you are paying for them to forget.

By contrast, when learning occurs close to the moment of need and is spaced and reinforced over time, retention and transfer rise sharply. Modern research shows that spacing learning into multiple touchpoints and tying it to real work outcomes makes behavior change several times more likely than one-off events.

> ***What*** we teach matters, but ***when*** we teach may matter just as much.

The Reality Check: Perfect Timing Is Impossible

Here's the part we usually skip: it is *not* realistic to perfectly time training around every future challenge. You cannot predict the exact day someone will encounter their first performance crisis, when a deal will turn political, or when a safety issue will become real.

You also can't run your entire L&D function as an emergency response unit. Some disciplines—safety, compliance, core technical knowledge—require foundations. Some skills must be practiced before they're tested in the wild. This is why organizations retreat to the calendar—budget cycles are predictable, vendor availability is predictable, rooms and Zoom slots are predictable. Real work is not.

The problem isn't that timing is hard. The problem is that we rarely design for timing at all. We accept a crude compromise—"*everyone in Q2*"—and then blame learners when nothing sticks.

The goal is not perfection. It's to ask a different question: *If timing is a powerful lever, how much closer to the moment of need could we get with the tools and data we already have?*

How People Actually Learn at Work

Look at how employees really learn—the misalignment becomes obvious.

Recent analyses of workplace learning show repeating patterns: most employees prefer to learn **while they are working**, not in abstract offsites. Many want to learn **at their own pace**, in short bursts. A large share say they only want formal training **when it's truly necessary**—when a real task or problem forces the issue.

Yet employees still receive most of their formal learning in **large, infrequent chunks**—multi-hour sessions or long modules delivered a few times a year.

So they do the rational thing: when they get stuck, they search Google, YouTube, or Stack Overflow. They ping colleagues, scroll internal chats, or scan wikis. Increasingly, they turn to AI copilots for immediate help.

Quietly, employees have built a just-in-time learning ecosystem for themselves. Meanwhile, the organization congratulates itself on completions for programs nobody remembers.

The Hidden Costs of Getting Timing Wrong

Bad timing doesn't show up on the dashboard, but it shows up everywhere else.

1. **Wasted training spend**
   When someone attends training months before they need the skill, decay does its work. By the time application finally arrives, most of the investment has evaporated. Organizations then pay again—through refresher courses, coaching, or avoidable mistakes.

2. **Performance gaps during the "wait"**
   Between the moment a new challenge appears and the next scheduled program, people improvise. Sometimes that's healthy; often it is expensive—in rework, escalations, lost deals, or safety incidents.
3. **Time pulled away from meaningful work**
   Long workshops and e-learning marathons consume time that could have been spent solving real problems. When the content doesn't align with immediate challenges, that time becomes a sunk cost, not an investment.
4. **Learned dependency**
   Over-reliance on scheduled training teaches people to wait: wait for the course before trying the tool, wait for the program before having the conversation. Over time, this erodes self-directed learning and reinforces the idea that "real learning" only happens when L&D schedules it.

Where Organizations Are Getting Timing (More) Right

Some organizations are quietly experimenting with timing in smarter ways.

In-app guidance instead of classroom training — Companies rolling out tools like Salesforce use digital adoption platforms to embed guidance directly into the system. Rather than asking people to remember steps from a one-time class, help appears in the interface when they're trying to perform the task. Result: faster onboarding, fewer errors, better adoption—without extra "courses."

Learning moments inside collaboration tools — Enterprise platforms increasingly integrate into Microsoft Teams, Slack, and similar tools. Short videos, checklists, and scenarios surface when someone searches or clicks a help icon in the middle of real work.

Peer-powered help at the moment of need — Programs like Google's "Googler-to-Googler" (g2g) model make it easy for employees to tap colleagues for short, targeted learning sessions just when they need them—rather than waiting for the next formal offering.

None of these solve the timing paradox completely. But they represent a deliberate shift from training scheduled around topics to support aligned with **moments**.

A More Pragmatic Timing Strategy: Events, Signals, Choices

If perfect timing is impossible, what's the realistic alternative? Think in terms of three types of triggers.

1. **Event-based triggers** — moments you *can* predict: promotions or role transitions, reassignment to a new customer segment or region, rollout of a new system, product, or regulation.

   Instead of a generic "annual leadership program," align intensive support around the 90 days before and after a role change. Instead of a one-off three-hour tool training, design a short pre-launch primer plus in-app support for the first month of use. The timing isn't perfect, but it's far closer to reality than "same course for everyone in April."

2. **Signal-based triggers** — data points that suggest a need emerging right now: spikes in customer complaints around a particular issue, increased errors in a process or system, deals repeatedly stalling at the same sales stage, safety or quality metrics dipping in a specific site.

Modern systems already generate these signals. AI makes it easier to detect patterns and trigger targeted support: a short scenario, a checklist, a 10-minute clinic pushed to the teams experiencing the problem *this week*, not next quarter.

3. **Choice-based triggers** — moments when people self-declare need: searching the learning platform or knowledge base, asking "How do I…?" questions in chat tools, prompting AI assistants with specific tasks.

This is the purest form of good timing: the learner has felt the gap and is motivated to close it. If your ecosystem is designed so that these "help me" moments surface not just answers, but small practices, templates, and job aids, you're harnessing timing instead of fighting it.

What AI Actually Changes About Timing

AI doesn't magically fix timing, but it removes key constraints.

From courses to copilots — When AI is embedded inside the tools people use (CRM, office suites, code editors), help can appear at the exact moment of need. The system can answer "show me three ways to respond to this email" *inside* the email client, not two months earlier in a classroom.

From generic nudges to signal-driven prompts — Rather than blasting everyone with "learning reminders," AI can look at patterns of work and prompt learning where friction is highest: "You've had several tickets escalated for the same issue. Would you like a short scenario and checklist before your next call?"

From static paths to adaptive pacing — Traditional learning paths assume a fixed sequence and timeline. AI can adapt based on behavior: accelerating for those who demonstrate competence quickly, slowing down and adding practice for those who struggle—before they hit a high-stakes moment.

The technology is already here. What's missing in most organizations is a timing strategy that connects the pieces.

Where to Start Without Burning It All Down

You don't need to cancel your entire training calendar. But you do need to prove—to yourself and your stakeholders—that timing is a real lever, not a philosophical nicety.

Practical starting moves:

1. **Redesign one critical transition**
   Pick one high-stakes role step (e.g., new manager) and rebuild the experience around a 90-day window with short interventions, manager check-ins, and in-flow resources instead of a single big event.
2. **Instrument one important process**
   Choose a process where mistakes are costly. Use system data to see where people struggle and overlay just-in-time job aids, guidance, or AI prompts at those steps.
3. **Turn one content area into an AI-first experience**
   Take a topic like "difficult conversations" and design resources explicitly for use through AI tools: prompts, scenarios, and checklists that people pull *when they're preparing for a real conversation*, not weeks in advance.
4. **Track timing, not just completion**
   Begin capturing when people access resources relative to key events: before, during, or long after. Over time, you'll see in your own data that "closer to the moment" often means "more likely to stick."

The Real Question

The timing paradox exposes what corporate learning has really been optimized for: budget cycles, not brain cycles. Vendor availability, not learner readiness. Completions we can report, not capabilities people can use under pressure.

You will never control when life throws people into the deep end. But in a world where AI sits inside our daily tools, where data reveals where friction is rising, and where employees are already hacking their own just-in-time learning, pretending that timing is a minor detail is no longer credible.

The value of training has never resided in the workshop itself. It resides in the moment when what was learned meets a real problem that actually matters.

The question is not whether you can time everything perfectly. The question is: *How long are you willing to keep investing in learning experiences that arrive after the moment has passed—or long before it ever comes?*