



Whitepaper

Bring pain forward

Happy schedules, the gap, and the discipline of finding reality early

FTTM essentials series.

Three connected practices that good product organizations use to surface project reality early, while there is still time to act on it.

Prepared by

lateralworks
FTTM methodology

Date

May 2026
Synthesis paper

Online

lateralworks.com
FTTM essentials

Table of contents

Bring pain forward

Abstract	03
01 Bring pain forward	04
02 Drop the happy schedule	07
In their own words: the practitioner case	10
03 Know the gap	11
04 Before-the-fact behavior	14
References	17

Core thesis. Three practices share a single discipline: bring reality forward in time so the team can act on it. Accelerate the pain so urgency arrives early. Build a schedule honest enough to expose the gap. Track the trend so you see slippage before it lands. Each practice fails on its own without a host environment that rewards early honesty.

Overview

Abstract

Most development projects do not fail because the team lacked skill. They fail because the team did not see the failure coming until it was too late to act. The gap between target and reality already existed. Nobody named it. By the time the milestone review surfaced the slip, the cheap options were gone.

This paper describes three connected practices that good product organizations use to keep that from happening. Each practice has a short name and a long history of being ignored. Bring pain forward. Drop the happy schedule. Know the gap. Used together, they form the discipline lateralworks calls before-the-fact behavior.

The practices are visible. The behavior that produces them is not. Optimistic schedules persist because the people building them face real consequences for naming risk. Pain stays low at the start of a project because nobody wants to be the team member who killed the mood in the kickoff meeting. Gaps stay hidden because exposing one is read as a personal failure rather than a system signal. Each of these patterns has been described in research on optimism bias, planning forecasts, and organizational defensive routines [1, 2, 3].

The fix is not exhortation. Telling people to be more honest about their schedules does nothing if the manager rewards the team that shows green and punishes the team that shows red. The fix is a host environment that makes early honesty the cheaper option, combined with a refresh cadence that gives trend signals time to accumulate before they turn into slips [4]. This paper walks through each practice, the academic and field evidence behind it, and the host conditions that let it stick.

The paper is organized as four sections. Sections one through three present each practice with its diagram, its failure modes, and the published research that supports it. Section four pulls the three together into a single picture of before-the-fact behavior and the role the host plays in keeping it alive.

01

Practice one

Bring pain forward

At the start of every project, the pain level is low. The team is fresh. The schedule shows green. The target sits comfortably in the future. Everything looks great. The pain stays low for weeks, sometimes months. Then a milestone arrives, the schedule slips, and pain spikes in a band so narrow there is no time to do anything thoughtful about it.

This is the curve the diagram on page five shows in green: pain flat at the floor until the target, then a vertical spike when reality lands. The yellow curve is the same project run by a team that practiced the opposite. They accelerated the threshold of pain. They forced urgency to arrive early, while the cheap options were still on the table. The total pain over the project life is roughly the same. The distribution is what changes.

The behavior is simple to describe and counter-cultural to practice. Most management cultures reward calm during the early phase and punish concern. The team member who raises a risk in month one is asked to prove the risk. The team member who waits until the slip happens is asked to fix it. The incentive structure is exactly backwards for what the project actually needs [3].

Practice one

Why teams keep the pain flat

The flat-green curve is not laziness. It is a learned response to how organizations treat early bad news. Argyris named the underlying pattern skilled incompetence: people become very good at saying things that look responsible and produce results that are not [3]. Around schedules and risk, the pattern looks like this. The team leader assembles a plan from the deliverables senior management asked for. The deliverables imply a schedule. The schedule has gaps the leader can see. Naming the gaps requires arguing with the people who set the deliverables. So the gaps stay implicit, wrapped in language like "aggressive but achievable" or "stretch target." Everyone in the room knows the schedule is not real. Nobody says so.

Argyris called the rules that hold this pattern in place defensive routines. The first rule is that the gap is undiscussable. The second rule is that the fact that the gap is undiscussable is also undiscussable [3]. A team member who breaks rule one is read as difficult. A team member who breaks rule two is read as a threat. The cost of breaking either rule is borne by an individual. The cost of keeping the rules in place is borne by the project months later, when the slip becomes visible.

Lovullo and Kahneman described the same pattern from the cognitive side. Executives, they argued, suffer from delusions of success: a systematic tendency to ignore base rates, anchor on best-case scenarios, and dismiss historical comparison data even when it is available [5]. The cognitive bias and the organizational defense reinforce each other. The bias makes the optimistic plan feel plausible. The defense makes the realistic plan feel disloyal.

How high performers invert it

High-performing product teams do not eliminate the bias or the defense. They build a host environment that makes the cost of early honesty lower than the cost of late surprise. The mechanism is straightforward. Bring forward the conversation that would normally happen at the slip review. Have it in the kickoff. Have it again at every weekly refresh. Reward the team member who spots a future failure now. Treat the spot as a contribution, not a complaint.

lateralworks engagement data show the pattern across multiple industries [10]. Teams that converged on schedule realism within the first month of a program ran two to three times the rate of on-target delivery compared with teams that protected an aggressive plan into month four or five. The difference was not in the people, the technology, or the difficulty of the work. The difference was when the team got honest.

Figure 1

The two pain curves

The diagram below shows both curves on the same axis. The green curve is the head-in-the-sand pattern: pain hugs the floor until the target, then climbs vertically when the slip lands and there is no time left to do anything but spend, pray, or both. The yellow curve is the inverted version: pain accelerates at the start, peaks well before the target, and falls as the team works through identified risks with months of runway to do it. Both paths end at the target. Only one of them gets there.

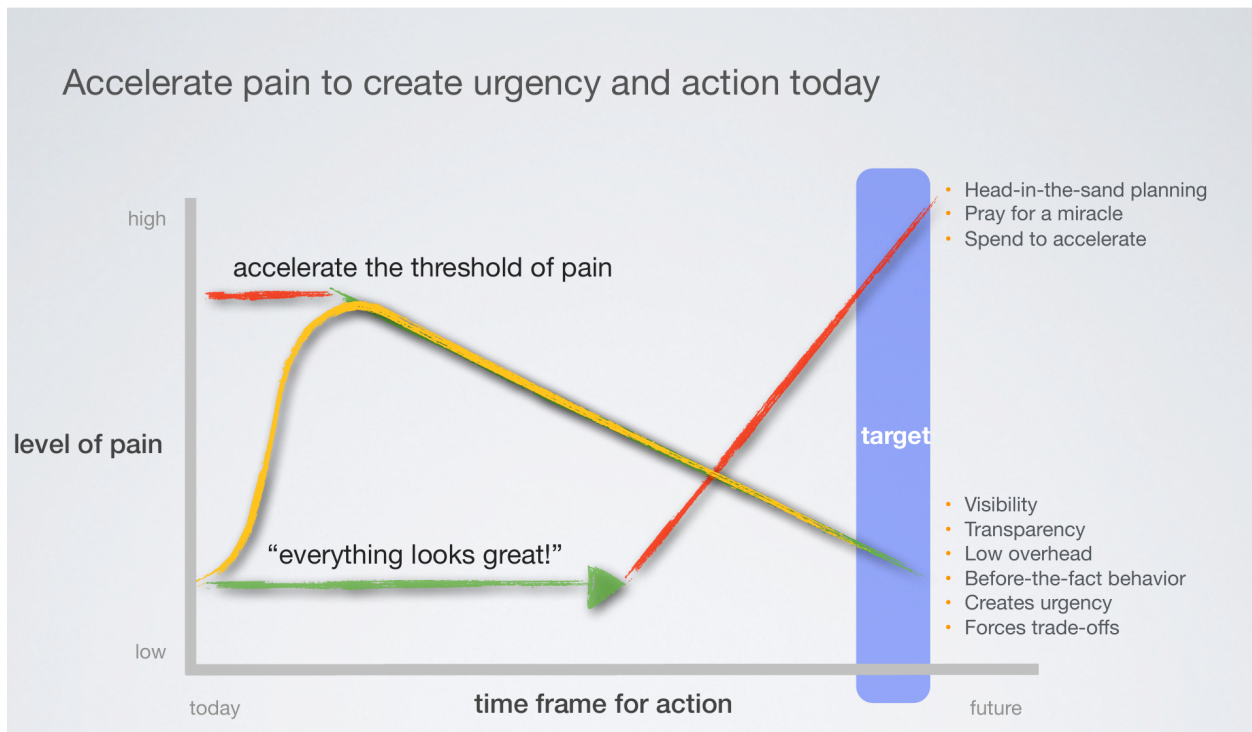


Figure 1. The accelerated pain curve versus the head-in-the-sand curve. Green: pain stays flat until the target, then spikes. Yellow: pain accelerates early, peaks before the target, then falls.

The right-side bullets on the diagram name the consequences of each path. The head-in-the-sand path produces three predictable late-stage responses: pretend the problem is not there, hope the team gets lucky, or throw money at acceleration once the slip is visible. The accelerated path produces visibility, transparency, low overhead, and the trade-off conversations that should have happened at the kickoff. The diagnostic question for a project team is which path the data show right now. If pain is rising in month nine, the project is on the green curve and the option set is already narrow.

02

Practice two

Drop the happy schedule

A happy schedule is the visible artifact of the head-in-the-sand pattern. It is the plan that finishes on time at the moment of kickoff because somebody wanted it to. Risk is absent from it. Learning cycles are absent from it. Right-first-time is the assumption. The schedule cannot be wrong because nothing in it has happened yet.

The empirical literature on happy schedules is large and consistent. Buehler, Griffin, and Ross documented the planning fallacy in a series of experiments showing that people systematically underestimate task completion times even when they have personal experience of similar tasks running long [1]. Flyvbjerg, working on infrastructure megaprojects, found cost and schedule overruns of roughly the same magnitude across decades, project types, and geographies, and attributed them to a combination of optimism bias and strategic misrepresentation [2]. The pattern repeats at every scale.

Practice two

What makes a schedule realistic

A realistic schedule does three things a happy schedule does not. It factors in risk as duration, not as a separate risk register that nobody opens. It factors in learning cycles for the parts of the work that nobody on the team has done before. And it shows critical paths visibly enough that the people on those paths know they are on them. Without any of those three, a schedule is a wish list with dates attached.

The first move toward realism is participatory. A schedule built top-down and handed to the team is the wrong artifact for managing work, regardless of how skilled the planner was. Best-practice product teams build the schedule with the people who will execute it, in a room together, with enough time to argue about durations and dependencies. lateralworks field data show that these teams typically spend roughly three times as long in planning as normal teams [10]. The investment looks expensive at the start. It pays back on every refresh thereafter because the team owns the plan and notices when reality starts to drift from it.

Why happy schedules persist

Happy schedules are easier to defend in a status review. They tell the audience what the audience wants to hear. They protect the team from being asked the question that started the whole pattern: why are you the one team that cannot hit this date when every other team manages to? The honest answer often involves the words "the other teams have happy schedules too and theirs will slip in month nine just like mine," but that answer is not safe to say out loud.

The deeper reason is structural. Sull, writing on execution gaps, found that the average manager believes commitments made in planning are commitments, while the average executive treats them as opening positions in a negotiation [6]. The schedule the team commits to is rarely the schedule the executive layer treats as binding. Teams learn this. They produce schedules that read as commitments and behave as forecasts. The artifact looks rigorous. The data behind it is wishful.

Flyvbjerg distinguished between optimism bias (the cognitive tendency to underestimate) and strategic misrepresentation (the deliberate underestimation of cost and time to secure project approval) [2]. He found that strategic misrepresentation was the larger of the two in public infrastructure. In private product development, the same combination shows up under different names. Sandbagging. Stretch goals. Aggressive baselines. The vocabulary is different. The behavior is the same.

Figure 2

Happy versus realistic

The diagram below contrasts the two schedule shapes. The happy schedule shows a comfortable runway followed by a delay block that nobody planned for. Urgency arrives only when the team realizes the runway has run out. The realistic schedule shows the same total work but acknowledges the risk and learning cycles as part of the plan, which leaves room for pull-in actions to bring the finish date in toward the target rather than push it past.

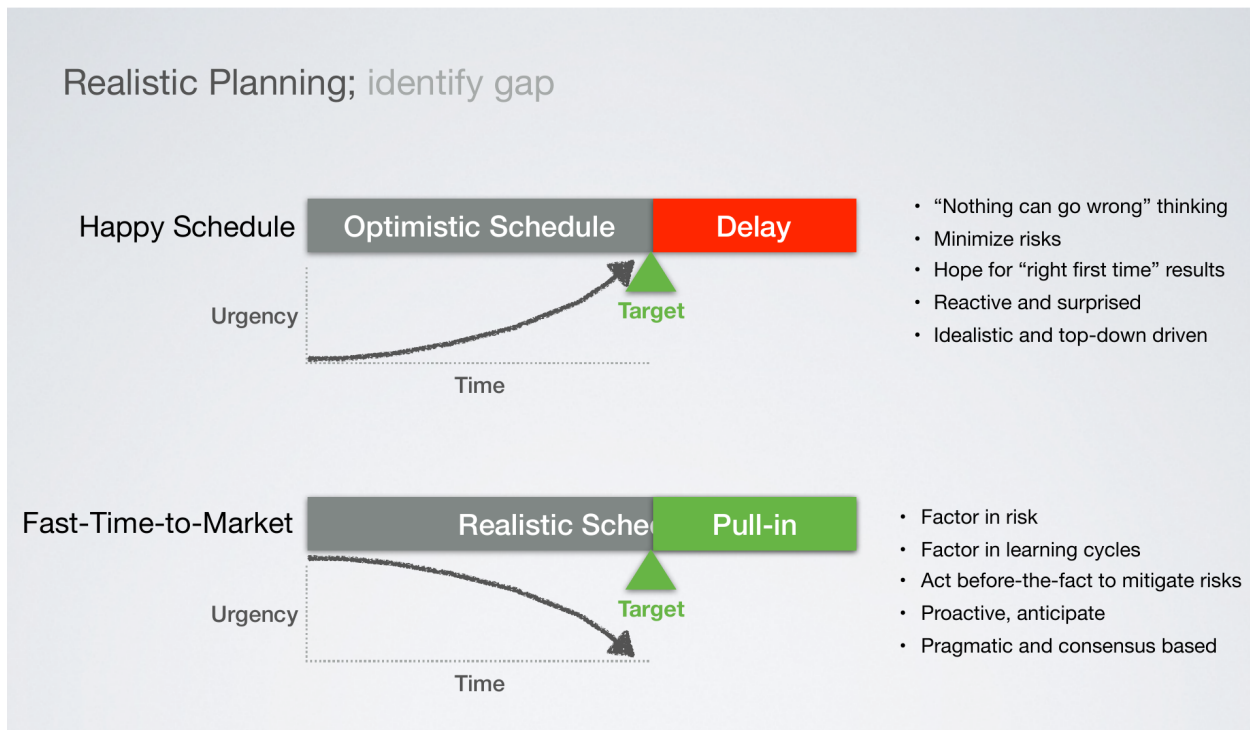


Figure 2. Happy schedule versus fast-time-to-market schedule. The happy schedule discovers its delay after the fact. The realistic schedule plans for risk and pulls in toward the target.

The bullets on the right capture the mindset behind each shape. The happy-schedule mindset assumes nothing will go wrong, minimizes risk in the plan, hopes for right-first-time results, reacts to surprises, and treats the schedule as a top-down mandate. The fast-time-to-market mindset factors risk explicitly, builds in learning cycles, acts before the fact to mitigate, anticipates instead of reacts, and treats the schedule as a consensus artifact owned by the people doing the work. The second mindset finishes more projects on time. The first mindset reports more projects on track at month three.

In their own words

The practitioner case

**The question is not
whether a gap exists.
It almost always does.
The question is whether
you want to know it,
or discover it too late.**

Neal Mitchell

lateralworks, The scheduling debate

03

Practice three Know the gap

The gap is the distance between the date the team is committed to and the date the team is on track to actually hit. The gap exists on almost every project. The variable is not whether there is a gap. The variable is whether anyone outside the project sees it before the milestone review.

A team that knows the gap can act on it. They can negotiate scope. They can add resource. They can compress critical paths. They can rebaseline the target. None of those options exist if the gap is invisible. The diagram on page twelve shows what visibility looks like in practice: a wigglechart that tracks the predicted finish date of a milestone refresh by refresh, week by week, so the trend is legible long before the slip lands.

Practice three

Why teams hide the gap

Teams hide the gap for the same reasons they keep the pain flat and write happy schedules. The cost of naming a gap is paid immediately by the person who names it. The cost of leaving it hidden is paid later by the project. Edmondson named the condition that makes early honesty rational: psychological safety, the shared belief that the team is safe for interpersonal risk-taking [7]. Without it, the rational move for any individual is to say less.

A second mechanism reinforces the silence. When schedules are tracked only by what is reported at status reviews, the only data points the organization sees are the ones the team chose to surface. The intermediate signal is missing. A milestone that has been quietly slipping by three days a week for two months looks healthy until the cumulative slip is too large to hide. The gap was always there. The system was not configured to show it.

The trend is the signal

The discipline that exposes the gap is trend tracking. Every refresh, the team records the predicted finish date for each target milestone. Over weeks, those dates form a line. The line is what the wigglechart plots. A flat line means the team is holding the milestone. A line that drifts up means the milestone is slipping. A line that drifts down means it is pulling in. The intercept of the trend with the target date tells the team how long they have before they need to act.

The analogy to financial markets is exact. A single day of a stock price says nothing. The trend of the stock over months says something about the health of the company. A single weekly schedule update says nothing. The trend over six weeks says whether the project is on the way to its target or moving away from it. Most project organizations track only the latest update. They have the daily prices and not the chart.

Reinertsen made the related point that the cost of delay on a product development program is rarely calculated and almost never used in trade-off decisions [8]. Teams will spend weeks choosing between two technical options without knowing what a week of delay costs. The wigglechart does not solve that directly. It does make the delay visible enough that the cost conversation can happen on the basis of data.

Figure 3

Trend analysis in practice

The diagram below shows a single milestone tracked over four months. The status header at the top reports a twelve-day slip against a 9-October target. The chart below it shows the trend that produced that slip. The predicted finish date drifted upward from late September to mid-October over a span of weeks. The slope of the trend line projects forward to an intercept with the target date. The red star marks the point at which the trend made the slip mathematically certain. Anyone watching the chart at that point had information to act on.

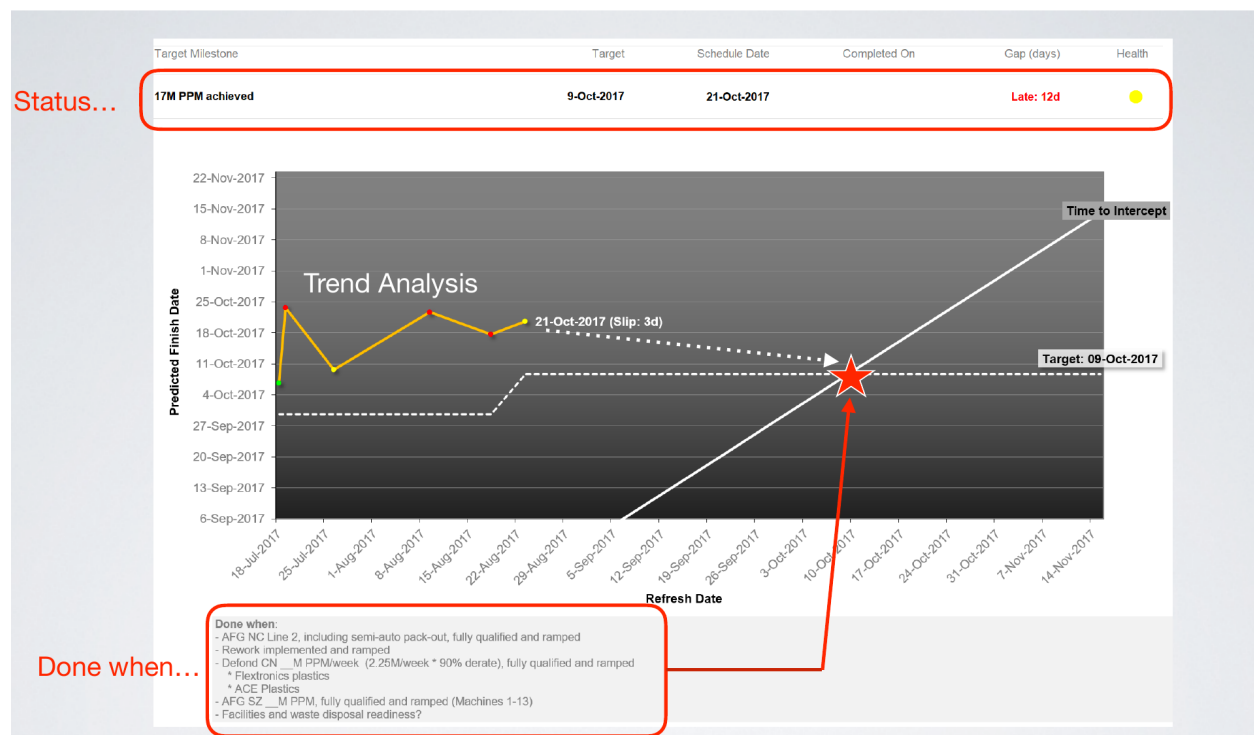


Figure 3. A wiggglechart for a single target milestone. The yellow line plots the predicted finish date refresh by refresh. The dotted extrapolation projects the trend forward. The red star marks where the slip became inevitable.

The "done when" box at the bottom of the diagram is the second piece of the practice. A trend chart says when the milestone will land. A done-when statement says what landing actually means: the list of specific deliverables that must be true for the milestone to be called complete. Without an unambiguous done-when, the milestone moves around even when the schedule stays still. Teams declare success on a partial set of deliverables and discover later that the missing pieces were on the critical path of the next milestone.

04

Synthesis

Before-the-fact behavior

The three practices share a single underlying behavior. Bring pain forward moves urgency from the end of the project to the beginning. Drop the happy schedule replaces wishful planning with a plan honest enough to expose the gap. Know the gap tracks the trend that tells the team how much time they have to act. All three are instances of acting before the fact instead of after it.

Before-the-fact behavior is uncommon because the host environment in most organizations rewards after-the-fact behavior. The team that finishes on time, however many weekends were burned, gets the promotion. The team that flagged the risk in month one and rebaselined to a realistic date gets the question about whether they were really committed to the original target. Until the host environment changes, no amount of training in the three practices will produce sustained behavior change.

Synthesis

What the host has to provide

The three practices need four host conditions to take root. lateralworks observes the same conditions repeatedly across client programs, regardless of industry [10]. Without them the practices erode within a year of a transformation, regardless of how cleanly they were installed.

First, the host has to ask the right question at the right time. The first question on a new program is not "how confident are you?" The first question is "what is the gap between your realistic finish date and the target, and what would it take to close it?" The first question signals what the host actually wants to know. If the first question rewards confidence, the team will produce confidence. If the first question rewards a specific gap number, the team will produce a gap number.

Second, the host has to refresh the plan at a cadence that matches the rate of learning. Weekly refresh is the lateralworks default. The point of the cadence is not to manage by metric. The point is to keep the trend chart populated densely enough that the slope is legible. Monthly refresh is too slow on programs where the learning rate is fast. Daily refresh is too noisy on programs where the learning rate is slow. The right cadence is the slowest one that still produces a readable trend [11].

Third, the host has to make problem identification cheaper than problem concealment. The team member who walks into the weekly refresh with a slip to report should leave the meeting with a pull-in plan and a sense of having contributed. The team member who reports green every week, on a program where green is mathematically improbable, should be asked what is being missed. The asymmetry of incentives is the lever. Without it, the practices revert to performance.

Fourth, the host has to distinguish targets from schedules. A target is what the business needs. A schedule is what the team can do. The gap between them is the actionable information. Treating the target as the schedule collapses the distinction and produces happy schedules by definition. The team has no way to report a gap because the schedule is required to equal the target. Reinstating the distinction is often the single highest-leverage change a host can make.

Where to start

For a leader who recognizes the patterns in this paper on a current program, the first move is the wigglechart. Install trend tracking on the three or four most important target milestones. Run it weekly for six weeks. The chart will surface reality on its own. Once the trend is visible, the gap conversation becomes possible on data rather than on opinion. The other two practices follow more easily once the gap is something everyone can see.

For a leader rebuilding the operating system rather than rescuing a single program, the sequence is reversed. Start with the host conditions. Set the first-question habit. Set the refresh cadence. Set the incentive asymmetry. Set the target-versus-schedule distinction. With those in place, the three practices propagate through the project portfolio with roughly the speed at which leaders begin asking the new questions in status reviews.

What changes when it is working

A program running on before-the-fact behavior produces a different set of artifacts. Status reviews open with the trend chart rather than the percent-complete number. Risk registers have specific owners and specific due dates rather than descriptive prose. Slip events are followed by pull-in actions within the same week, not by escalation memos two weeks later. The team can name its current critical path and its second critical path without consulting a document. The host can state the gap on every active program from memory.

None of those artifacts are difficult to produce. They are unusual because the host environment that produces them is unusual. The three practices in this paper are the symptom. The host is the cause. A program that wants the symptom without the cause will get neither for long.

Diagnostic prompt. Pick a current program. Ask the team leader two questions. What is the gap between your realistic finish date and the committed target? Where is the trend over the last six weeks? If the answers are confident and specific, the host is working. If the answers are evasive, the host is the work.

Sources

References

- [1] Buehler, R., Griffin, D., and Ross, M. "Exploring the planning fallacy: Why people underestimate their task completion times." *Journal of Personality and Social Psychology*, 67(3), 1994, pp. 366-381.
- [2] Flyvbjerg, B. "Curbing optimism bias and strategic misrepresentation in planning: Reference class forecasting in practice." *European Planning Studies*, 16(1), 2008, pp. 3-21.
- [3] Argyris, C. "Skilled incompetence." *Harvard Business Review*, September-October 1986, pp. 74-79.
- [4] lateralworks. "The rhythm of accountability." Ideas archive, 2019.
<https://lateralworks.com/ideas/2019/1/20/the-rhythm-of-accountability>
- [5] Lovallo, D., and Kahneman, D. "Delusions of success: How optimism undermines executives' decisions." *Harvard Business Review*, July 2003, pp. 56-63.
- [6] Sull, D., and Spinosa, C. "Promise-based management: The essence of execution." *Harvard Business Review*, April 2007.
- [7] Edmondson, A. "Psychological safety and learning behavior in work teams." *Administrative Science Quarterly*, 44(2), 1999, pp. 350-383.
- [8] Reinertsen, D. G. *The Principles of Product Development Flow: Second Generation Lean Product Development*. Celeritas Publishing, 2009.
- [9] Kahneman, D., and Tversky, A. "Intuitive prediction: Biases and corrective procedures." In Kahneman, D., Slovic, P., and Tversky, A. (eds.), *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, 1982, pp. 414-421.
- [10] lateralworks. Internal assessment database. Engagement data across client programs, 2015-2026.
- [11] lateralworks. "Pull-in before you slip." Ideas archive, 2019.
<https://lateralworks.com/ideas/2019/6/23/pull-in-before-you-slip>
- [12] lateralworks. "Accelerate pain." Ideas archive, 2019. <https://lateralworks.com/ideas/2019/1/17/accelerate-pain>
- [13] lateralworks. "Obvious, but not really." Ideas archive, 2019.
<https://lateralworks.com/ideas/2019/10/16/obvious-but-not-really>
- [14] Argyris, C. *Overcoming Organizational Defenses: Facilitating Organizational Learning*. Prentice Hall, 1990.
- [15] lateralworks. "It's the second or third critical path that can sneak up and bite you." Ideas archive, 2024.
<https://lateralworks.com/ideas/2024/12/14/its-the-second-or-third-critical-path-that-can-sneak-up-and-bite-you>
- [16] lateralworks. *FTTM Self-Paced Tutorial and Reference Guide*, Section on wiggglechart trend analysis, 2019.
- [17] Kotter, J. P. "Leading change: Why transformation efforts fail." *Harvard Business Review*, March-April 1995, pp. 59-67.
- [18] Goldratt, E. M. *Critical Chain*. North River Press, 1997.
- [19] Buehler, R., Griffin, D., and Peetz, J. "The planning fallacy: Cognitive, motivational, and social origins." *Advances in Experimental Social Psychology*, 43, 2010, pp. 1-62.
- [20] lateralworks. "The scheduling debate." Working paper, 2025.