

NOTES FROM THE MARGINS

20 ideas from the books we read, each with the reason it holds and the place it stops holding.

YOU CANNOT TRUST ANYONE'S REPORT OF WHAT THEY USED TO BELIEVE, INCLUDING YOUR OWN, BECAUSE A CHANGED MIND OVERWRITES ITS OWN PRIOR STATE RATHER THAN STORING IT ALONGSIDE THE NEW ONE.

In the experiments described, attitudes are measured, a persuasive message is delivered, attitudes shift measurably — and then participants asked to recall their original position report something close to their new one. The recall task substitutes the currently available belief for the unavailable former one, and people often deny they ever thought otherwise. So the very people who changed most are the least able to testify that change occurred.

Where it stops holding: The effect is demonstrated on soft attitudes like the death penalty; beliefs tied to a vivid, dated event ("I was against this until the outage in March") may be far more recoverable.

So: Never evaluate whether something persuaded people by asking them afterward — capture the baseline before exposure, in writing, and diff against it; retrospective "did this change your mind?" surveys systematically report zero effect where a real effect exists.

From Thinking, Fast and Slow by Daniel Kahneman

**IN MOST PURCHASES SOMEONE MAKES ON BEHALF OF
AN ORGANIZATION, THE BUYER IS NOT OPTIMIZING
FOR THE BEST PRODUCT BUT FOR THE LEAST
BLAMEABLE CHOICE, SO MARKET LEADERSHIP IS
ITSELF A FEATURE COMPETITORS CANNOT MATCH
WITH QUALITY.**

Fisher observes that corporate staff recommended IBM over rivals whose machines they privately judged better or cheaper, because a failure with the leader costs them nothing while a failure with an unknown vendor costs them their neck. The decision-maker's downside is personal and asymmetric, while the upside of picking better accrues to the employer. Reputation therefore functions as insurance the recommender buys with the company's money.

Where it stops holding: This asymmetry weakens where the buyer is also the owner, or where procurement is judged on measured outcomes rather than on absence of complaints, which is increasingly common in bottoms-up software adoption.

So: Stop selling on feature superiority to enterprise buyers and start manufacturing blame cover — named reference customers in their industry, contractual guarantees, third-party audits — so a champion who picks you can survive your failure.

From Common Stocks and Uncommon Profits and Other Writings by
Philip A. Fisher

**TWO PEOPLE WITH IDENTICAL WORKLOADS AND
COMPARABLE WORK CAN DIFFER TWOFOLD IN
THROUGHPUT PURELY BY THE ORDER THEY WORK IN,
SO SEQUENCING IS A PRODUCTIVITY LEVER FULLY
INDEPENDENT OF EFFORT OR TALENT.**

In the Milan study, cases were randomly assigned and the workload was equal, which strips out both selection and difficulty as explanations; the only surviving difference was whether a judge ran cases concurrently or in sequence, and that alone separated 178 days from 398. A 50% increase in multitasking bought roughly a 20% increase in days to completion. Nobody worked less hard and nobody had easier cases — the slow judges paid for their own scheduling policy.

Where it stops holding: Judges who held many cases open may have been doing something the study can't see — waiting on external parties, or deliberately parallelising to hedge against blocked cases — so the causal arrow could partly run from case circumstances to apparent multitasking.

So: Before adding people or hours to a backlog, cap work-in-progress and finish items in sequence; the random assignment in this study means you can expect the gain from queue discipline alone, without hiring or working longer.

From Great at Work by Morten Hansen

AN INDUCEMENT'S JOB IS FINISHED THE MOMENT A DECISION IS MADE, BECAUSE PEOPLE IMMEDIATELY MANUFACTURE FRESH REASONS FOR THE CHOICE — SO THE INDUCEMENT CAN BE WITHDRAWN AND THE COMMITMENT WILL STAND ON REASONS THE PERSON INVENTED THEMSELVES.

Cialdini's low-ball works because after committing, a buyer generates new justifications that were not part of the original offer; removing the price that caused the decision leaves those new struts holding it up. The commitment therefore survives the disappearance of its own cause, and the person experiences the decision as self-authored rather than purchased.

Where it stops holding: If the withdrawal is blatant or the new reasons are thin, people do walk — the tactic depends on the inducement being removable in a way that reads as circumstance, not betrayal.

So: Stop treating launch discounts, free tiers and signing bonuses as recurring costs: price them as one-time commitment triggers, and design the post-commitment moment so users are asked to state their own reasons for being there before the inducement lapses.

From Influence by Robert B. Cialdini

AN EXPERT CANNOT RELIABLY TELL YOU WHAT IS HARD ABOUT THEIR FIELD, BECAUSE THE ACT OF GETTING GOOD COMPILES THE DIFFICULTY OUT OF MEMORY.

The passage reports a professor whose lecture failed because he 'couldn't remember which concepts were hard' — expertise converts deliberate high-level reasoning into low-level script-like processes that leave almost no memory trace. The forgetting is not carelessness; it is a structural byproduct of the same compression that makes the expert fast.

Where it stops holding: Some experts do retain vivid memories of their own early confusion, and good teachers deliberately rehearse the novice view, so the amnesia is a tendency rather than a law.

So: Never source your onboarding docs, API ergonomics, or curriculum difficulty ratings from the most senior person; instrument where actual beginners stall, or make people write down what confused them while they are still confused.

From The Emotion Machine by Marvin Minsky

MULTITASKING DOESN'T JUST SPLIT YOUR TIME — IT FORCES YOU TO REDO THINKING YOU HAD ALREADY FINISHED, SO THE COST IS RE-DERIVATION, NOT DIVISION.

Since conscious thoughts happen one at a time, interleaving two tasks cannot save any steps; the total number of steps is fixed. But each return to an abandoned stream requires reconstructing the last conclusions, so steps get repeated — the interleaved trace contains the same sub-steps multiple times. The cost is therefore superlinear in the number of switches, not a wash.

Where it stops holding: For tasks with cheap external state — a checklist, a compiler error, a written intermediate result — the reload is nearly free, so the re-derivation penalty is a property of tasks held in the head, not of switching itself.

So: Stop optimizing for fewer interruptions and start optimizing for cheap resumption: end every work session by writing down the last conclusion reached, so the re-entry cost is a read instead of a rethink.

From Mental Traps: The Overthinker's Guide to a Happier Life by Andre Kukla

THE DAMAGE FROM A TASK SWITCH LANDS ON THE TASK YOU SWITCH *TO*, NOT THE ONE YOU LEFT — AND IT PERSISTS AFTER YOU HAVE FULLY MOVED ON.

Leroy's experiments forced switches mid-puzzle and then measured performance on an unrelated second task, finding degraded performance on the *new* task in proportion to how much residue the first one left. Crucially, residue appeared even when the first task was completed before the switch, so 'finishing' does not clear it — and it was thickest when the prior work was unbounded and low-intensity.

Where it stops holding: Lab-forced switches between artificial tasks may exaggerate residue relative to real work, where switches are often self-timed at natural breakpoints and the two tasks may share context.

So: Schedule the hardest task of the day first, before anything ambient has run — and prefer one bounded, intense session over a long shallow one, since low-intensity unfinished work leaves the thickest residue behind.

From Deep Work by Cal Newport

HABITUAL MULTITASKING RETRAINS WHAT YOUR ATTENTION TREATS AS SALIENT, SO THE DEFICIT PERSISTS INTO MOMENTS WHEN YOU ARE DOING ONLY ONE THING.

In the arrow-counting study, heavy multitaskers were not slower because they were multitasking during the test — they were doing a single task. They failed as distractor count rose because their attention had been trained to register every stimulus regardless of relevance. The cost is thus a durable change in filtering, not a momentary tax paid only while juggling.

Where it stops holding: This is correlational — people with weaker baseline filtering may simply gravitate toward multitasking, which would reverse the causal arrow entirely.

So: Treat single-task practice as training rather than as productivity hygiene: the payoff is measured in what you can ignore weeks later, so protect focus blocks even on days when nothing urgent is competing for them.

From Ikigai by Héctor García & Francesc Miralles

AN AUTOMATION IS FINISHED NOT WHEN IT RUNS, BUT WHEN A LESS-SKILLED PERSON CAN OWN IT — DELEGABILITY, NOT EXECUTION, IS THE COMPLETION TEST.

Limoncelli reports a colleague's litmus test: a backup system was 'automated enough' only once the company secretary could run the daily tape changes, with the system emailing status and instructions and telling her when to keep hands off. That standard forces you to encode the failure handling and the judgement calls, not just the happy path — which is why the system eventually ran for months without expert intervention.

Where it stops holding: For a task you will always own yourself, building the delegable version costs far more than the script that merely works, and the extra hardening may never pay back.

So: Stop grading your scripts by whether they complete; grade them by whether you could hand them to the least-skilled person on the team today, and treat every case where you'd say 'well, except when...' as the actual remaining work.

From Time Management for System Administrators by Thomas A. Limoncelli

WHEN A DOMAIN RESISTS AUTOMATION BECAUSE EVERY JOB IS ONE-OF-A-KIND, STANDARDIZING THE PRODUCT IS THE TEMPTING FIX AND THE ONE THAT FAILS — THE FLEXIBILITY HAS TO GO INTO THE MACHINE INSTEAD.

Construction has 'job lots of one', so fifty years of effort went into standardizing buildings to manufacture repetition; those attempts largely failed both on process rigidity and on customers refusing the standardized product. The passage's alternative is flexible robotics that absorb variety, built as the industry's own equipment rather than factory robots relocated to the site.

Where it stops holding: Plenty of industries did automate precisely by standardizing their output — shipping containers, fast food, prefab components — so the failure may be specific to buildings and their buyers rather than a general law.

So: Before forcing your inputs into a uniform shape so a pipeline can chew them, check whether the variety is what your users are actually paying for; if it is, spend the effort on tooling that tolerates variance and expect to build it around your existing workflow rather than importing someone else's.

From Springer Handbook of Robotics by Bruno Siciliano & Oussama Khatib

ASKING SOMEONE ABOUT THEIR INTENTIONS IS NOT A NEUTRAL MEASUREMENT — IT AMPLIFIES WHATEVER THEY ALREADY FEEL, MAKING THEM LIKELIER TO DO THINGS THEY VIEW POSITIVELY AND LESS LIKELY TO DO THINGS THEY VIEW NEGATIVELY.

Merely being asked whether you plan to buy a computer in the next six months raises the odds you buy one by 18 percent; the same questioning about flossing and avoiding fatty foods raises compliance. But asked about an unappealing act, thinking it through makes you less likely to do it. The question does not supply motivation; it forces the person to consult a stored evaluation and then commits them to it.

Where it stops holding: The effect sizes come from consumer and health studies where the behavior was already low-cost and near-decided, so the lever may be trivially small anywhere real friction exists.

So: Treat your surveys, onboarding questionnaires, and check-in prompts as interventions, not instruments: ask about the behaviors your users already feel good about, and stop asking prospective questions about the ones they find distasteful — you are actively talking them out of it.

From Give and Take: A Revolutionary Approach to Success by Adam M. Grant Ph.d.

EXPERTISE ERASES ITS OWN LEARNING HISTORY: AS SKILLS COMPILE INTO AUTOMATIC ROUTINES, THEY STOP LEAVING MEMORY TRACES, SO AN EXPERT LITERALLY CANNOT RECALL WHICH PARTS WERE HARD.

The passage describes a professor whose lecture failed because he could no longer remember which concepts were difficult — his high-level skills had been converted into low-level script-like processes that leave almost no trace in memory. The deficit is not arrogance or empathy failure; it is a storage side effect of the compilation that made him fast. That means the expert's introspective report about difficulty is missing data, not biased data.

Where it stops holding: Experts who teach constantly re-encounter novice difficulty and can rebuild an accurate model of it, so the amnesia is repairable rather than structural.

So: Stop sourcing your onboarding docs, tutorials and API ergonomics from expert introspection; capture where actual novices stall (recordings, first-run telemetry, support logs) and treat that trace as the ground truth the expert can no longer produce.

From The Emotion Machine by Marvin Minsky

DAILY USE OF A SKILL DOES NOT MAINTAIN IT — SKILLS THAT ARE USED BUT NEVER CONSCIOUSLY PRACTISED DECAY ANYWAY, SO 'USE IT OR LOSE IT' IS FALSE COMFORT FOR PROFESSIONALS.

The passage separates two things people conflate: exercising a skill in production and deliberately working on it. Once a skill is automatic, using it recruits the cached version rather than refining it, so performance can slide slowly in exactly the abilities someone exercises every day. The professional's belief that hard-won skills are now self-sustaining is what makes the decay invisible.

Where it stops holding: For many well-bounded skills, frequent real use with real feedback does hold performance steady, and the claimed silent decay is asserted here more than measured.

So: Audit your most-used core skills, not your rusty ones: schedule explicit practice or calibration against a standard for the things you do daily, since those are the ones no one will flag as deteriorating.

From *Badass: Making Users Awesome* by Kathy Sierra

TALKING TO CUSTOMERS BEFORE YOU HAVE A PRODUCT IS FOR FALSIFYING THE FOUNDERS' HYPOTHESES, NOT FOR COLLECTING THE FEATURE LIST — THE SPEC MUST STILL COME FROM THE VISION.

The passage explicitly forbids making a list of all features customers want, handing Product Development the sum of customer requests, or running focus groups; instead discovery turns founders' initial hypotheses into facts. The initial product specification comes from the founders' vision, and discovery is the search for customers who happen to share that problem. So the interview's output is a verdict on a guess, not a requirements document.

Where it stops holding: In a crowded, well-understood market the customers genuinely do know the missing features, and treating their requests as mere hypothesis-tests wastes free specification work.

So: Write down your business-model guesses as falsifiable statements before the first call, and score each conversation as pass/fail against them rather than logging feature requests into a backlog.

From The Startup Owner's Manual: The Step-by-Step Guide for Building a Great Company by Steve Blank & Bob Dorf

**A FOUNDER RETAINS THE RIGHT TO OVERRIDE
EVERY CUSTOMER HE INTERVIEWED, BUT ONLY
IF HE CAN ARTICULATE WHY THEY ARE WRONG
— THE LICENCE IS CONDITIONAL ON THE
ARGUMENT, NOT ON CONVICTION.**

The passage grants that in a new market a founder's vision of what can be may be clearer than potential customers', so ignoring the feedback is legitimate. But it attaches a condition: the founder must be able to state the 'why', not merely dismiss the data. That converts stubbornness from a personality trait into something auditable by other people.

Where it stops holding: Every founder who fails can also produce an articulate 'why', so the condition filters for eloquence rather than for correctness.

So: When you decide to override user research, write the override and its reasoning down at the time, so the claim can be checked later instead of being retrofitted after the outcome is known.

From The Startup Owner's Manual: The Step-by-Step Guide for Building a Great Company by Steve Blank & Bob Dorf

FLOW DOES NOT REDUCE THE EFFORT A TASK TAKES — IT REMOVES THE SECOND, HIDDEN EFFORT OF POLICING YOUR OWN ATTENTION, WHICH IS THE ONE THAT ACTUALLY DEPLETES YOU.

The passage separates two costs that usually get lumped together: concentration on the task, and the deliberate control of attention that keeps you on it. Riding a motorcycle at 150 mph or playing tournament chess is enormously effortful, yet in flow no self-control is spent holding attention in place, so the whole budget goes to the work. That is why the same person can grind for hours without willpower on one task and raid the refrigerator every ten minutes on another.

Where it stops holding: The two efforts may not be as separable as the account implies — sustained high-intensity work in flow still leaves people exhausted, which suggests a shared resource is being drained either way.

So: Stop trying to make hard work easier and start auditing what forces you to re-decide to keep working — notifications, ambiguous next steps, unclear whether the task is worth doing. Removing the re-deciding, not the difficulty, is what buys you long sessions.

From Thinking, Fast and Slow by Daniel Kahneman

CONCENTRATION HAS AN ACTIVATION-ENERGY SHAPE: ENTERING IT IS EXPENSIVE AND STAYING IN IT IS CHEAP, SO AN INTERRUPTION COSTS YOU THE ENTRY PRICE AGAIN, NOT THE MINUTES IT TOOK.

The passage draws the analogy to metastable physical systems that need external energy added up front to hold a higher-energy state they would otherwise decay out of. Effort is required to focus attention into flow, but once inside, external distractions are much less likely to disrupt it and even heavy exertion feels effortless. The default without that up-front expenditure is not neutral rest but decay into random, low-energy states.

Where it stops holding: The physics analogy is explicitly flagged as possibly superficial, and it says nothing about how large the entry cost actually is — for some tasks it may be trivial.

So: Price interruptions at the cost of re-entry, not the length of the interruption, and schedule accordingly: one three-hour block beats three one-hour blocks, and a five-minute Slack check during hour one is far more expensive than the same check during hour three.

From Creativity by Mihaly Csikszentmihalyi

SELF-CONSCIOUSNESS DURING WORK IS A CAPACITY PROBLEM, NOT A CHARACTER PROBLEM — THE SELF REAPPEARS WHENEVER THE TASK LEAVES ATTENTION UNSPENT.

The passage locates absorption in a specific condition: the activity consumes essentially all available attention, leaving no excess psychic energy to process anything else. The reported disappearance of the self — dancers, climbers and chess players all describing that they stop experiencing themselves as separate from the action — follows mechanically from that saturation rather than from any spiritual achievement. Comfortable, under-demanding work leaves spare capacity, and monitoring yourself is what fills it.

Where it stops holding: Saturating attention could equally produce anxiety or tunnel-visioned error rather than absorption, so full capacity use is clearly necessary but not sufficient.

So: When someone is stuck in self-critical rumination during work, raise the difficulty or tighten the feedback loop instead of coaching confidence — give the spare attention something to do. Design tasks to sit at the edge of capability, and treat easy work as the setting where

morale problems appear.

THE COST OF SWITCHING TASKS SCALES WITH THE COMPLEXITY OF THE TASKS, WHICH MEANS EVERYONE'S INTUITION ABOUT MULTITASKING IS CALIBRATED ON THE WRONG DATA.

Meyer's estimate puts the overhead at 25% or less for simple tasks but over 100% for complicated ones — a fourfold-plus difference in penalty depending on what you're switching between. People form their beliefs about multitasking from the cheap end (folding clothes while watching TV, glancing at a phone in a queue), where it genuinely costs almost nothing, then generalise that felt cheapness to spreadsheet-and-interruption work where the same behaviour more than doubles the time. The error isn't that people underestimate a constant cost; it's that they extrapolate from the regime where the cost is real but negligible.

Where it stops holding: Meyer's ranges come from lab task-switching paradigms and may not transfer cleanly to hours-long knowledge work, where recovery is partly a matter of external notes and tooling rather than raw cognitive reconfiguration.

So: Stop applying one interruption policy across the board: let simple work be interruptible and protect only the genuinely complex work,

because that's where the entire penalty lives — and buy tooling (checkpoints, written state, resumable context) whose job is to lower the restart cost of the hard tasks specifically.

REDUNDANCY BOUGHT OUT OF A FIXED ATTENTION BUDGET IS SELF-DEFEATING: EACH BACKUP YOU ADD MAKES EVERY OPTION WORSE, SO THE INSURANCE DEGRADES EXACTLY THE THING IT WAS SUPPOSED TO GUARANTEE.

Scott's five transport methods looked like hedging — if the motor sledges fail, use the dogs; if the dogs fail, the ponies — but the attention needed to source and master each was drawn from the same finite pool, and he ended up with second-rate ponies. On top of that, the options had to interoperate: they moved at different speeds, so the portfolio created a coordination problem that none of the individual choices had. Redundancy in a system with independent budgets per component genuinely adds reliability; redundancy in a system where all components draw from one attention budget converts one strong option into several weak ones.

Where it stops holding: For genuinely uncertain, high-variance bets, several mediocre independent shots really can beat one excellent shot — venture portfolios and evolutionary search both work this way — so the argument holds only where quality thresholds are sharp and options must interoperate.

So: When you catch yourself adding a fallback, first price it in attention taken from the primary and in the coupling it introduces, then ask

whether the fallback will still clear the minimum quality bar once it's been funded that cheaply — if not, delete it and put everything behind one path.