The Benefits of Failure

https://gerolds.github.io/textbook/textbook/posts/the-benefits-of-failure/

Contents

The Benefits of Failure

A résumé full of cancelled projects, broke studios, grant-dependent non-starters, and games nobody played looks like evidence of a person who cannot get things right. Sometimes it is. But sometimes it is the opposite: someone who has seen every way things go wrong, understood *why* they went wrong, and built a map of failure so detailed that the path to excellence becomes legible by exclusion.

This essay is about that second person, and about why the environments most people dismiss as career dead-ends can produce exceptionally capable practitioners, if the person inside them was paying attention.

---

1. The core claim

A career mostly spent in failing, mediocre, irrelevant, or misguided projects fails to teach one thing directly: what it feels like to ship something excellent while you are doing it. The daily texture of a well-run production, the rhythm of a team that trusts its own process, and the confidence of knowing the game works before the market confirms it are not available inside a studio that has never produced them.

But that career can still teach, to those who pay attention, something years inside a successful studio often leave implicit: **which conditions make excellence impossible, and therefore which conditions excellence probably depends on.** Not a complete recipe, and not a substitute for proximity to strong work, but a sharper understanding of the disciplines, refusals, and standards without which a team ships something people forget.

Thesis: Inside a failing organization, many variables that success normally hides become exposed. Process, taste, communication, scope, hiring, accountability, direction: all of it becomes easier to see because so little is working well enough to disappear into the background. A person who works inside that for years and *studies*

it develops a negative-space expertise. But negative-space expertise is only part of excellence. To become fully reliable, it has to be brought into contact with environments, peers, and artifacts that can provide higher-resolution feedback than failure usually can.

2. **Why dysfunction is legible**

In a well-run studio, success hides its own causes. The game ships well. The team felt good. Revenue came in. But which decisions actually mattered? Was it the creative director's taste, or the production discipline that kept scope honest? Was it the team's chemistry, or the hiring filter that assembled the team in the first place? Was it the game's design, or the market timing? Success is a compound signal. It is very difficult to decompose.

Failure, by contrast, is analytically transparent.

When a game ships and nobody cares, you can usually trace the cause. The core loop was never proven. The scope inflated past the team's capacity. Nobody had the authority to cut. The pitch was funded, not the game. Direction changed three times and nothing was thrown away. The team built systems instead of experiences. Feedback was never sought, or sought too late, or ignored when it arrived.

Each of these is a discrete, observable failure. You don't need a textbook to identify it. You just need to look at the gap between what the team intended and what it actually produced, and ask *why*.

In a successful environment, you experience the answer without ever seeing the question. In a failing environment, you experience the question every day. **The question is the education.**

3. **The laboratory of low stakes**

Failing organizations share a structural feature that is rarely appreciated: they are full of space. Not space in the sense of comfort or resources. Slack, autonomy, low scrutiny. Nobody is watching closely because nobody has built a process rigorous enough to watch closely. Accountability is vague. Ownership is

informal. There is no entrenched pipeline that breaks when someone does something differently.

This creates a laboratory. Not a well-funded one. A messy one, with no controls. But a person inside it can run small experiments that would be structurally impossible in a tighter organization.

Build a small game for an exhibition nobody important will attend. Make it in three weeks on no budget. Watch it become the only thing visitors aged 10–25 remember. Then look at the ten other things built by colleagues under the same constraints and notice they did not produce the same result. The variables are almost identical (timeline, budget, tools, audience, context), so the difference must live in the decisions. That difference is data.

This is not a controlled experiment. It is a paired comparison in the wild, where the person running it is also the one interpreting the result, and the feedback is immediate, embodied, and difficult to dismiss as someone else’s context.

Still, there is a limit. Low-stakes environments are good at revealing what obviously fails. They are often much worse at revealing what separates the merely competent from the genuinely excellent. The audience may be too forgiving, the peers too inexperienced, the incentives too weak, and the consequences too noisy. So the laboratory matters, but only if the person inside it keeps looking beyond it.

The same logic applies across every role. A programmer who builds three prototypes and notices that one of them is the only one the team actually plays has learned something about what makes interaction sticky. A designer who writes four pitch documents and notices that one of them is the only one that makes a funder lean forward has learned something about clarity. An artist who tries two visual directions and notices that one of them makes strangers stop scrolling has learned something about signal density.

None of these experiments require permission, budget, or institutional support. They require someone who treats a low-stakes environment as a classroom instead of a waiting room.

4. **What the successful environment fails to teach**

Most people assume the best way to learn excellence is to be inside it. Work at a great studio. Ship a great game. Absorb the culture.

This works, sometimes, for a specific kind of learning: pattern replication and standards acquisition. A person who spends five years at a studio with strong creative direction will absorb some of that direction's instincts. They will learn what "good" feels like in daily reviews, what the bar is for acceptable work, what gets praised and what gets reworked.

What they often will not learn is **why** those patterns exist, or what would break if the patterns changed, or how to rebuild them in a different context. The patterns are inherited, not derived. The person knows the answer without knowing the problem.

This is the profile of someone who may be excellent *inside the system that trained them* yet less transferable than they appear. They joined an organization where a few wise decisions, made years earlier by people who are sometimes no longer there, produced a self-reinforcing culture of quality. The culture carries more of the developmental burden than the individual realizes. When they leave and try to reproduce the result somewhere without that culture, they discover that what they thought was entirely personal capability was partly institutional momentum.

Success makes its own causes invisible. The person inside it learns what excellence looks like but not what generates it. They learn the posture, the vocabulary, and the taste, the outputs, without necessarily learning the inputs. When the inputs are needed, they reach for what they remember: the surface, not the engine.

The person who came up through failure has the opposite problem and the opposite advantage. They have never experienced the daily texture of excellence. They might not know the posture or have the vocabulary. But they may know, with unusually high resolution, what prevents excellence from emerging. They have watched it fail to emerge in dozens of configurations, and they have mapped the failure modes.

Given the right opportunity, this person can sometimes build. Not just replicate. **Build.** But only if their negative knowledge is eventually disciplined by contact with real standards: strong peers, demanding audiences, clear comparisons, and honest feedback from outside the failing system that trained them. Otherwise they risk becoming brilliantly diagnostic and only intermittently excellent.

5. **The eyes that see it**

Not everyone who works in a failing organization learns from it. Most don’t. Most adapt to the local standard, collect their pay, and develop increasingly sophisticated justifications for why the work doesn’t land. The organization provides cover. There is always a reason the game didn’t sell: the market, the budget, the timeline, the publisher, the platform. The reasons are usually partially true, which makes them effective shields.

The person who learns is the one who refuses the shield.

They look at the same failed project and ask different questions. Not “why did this fail despite our effort?” but “what about our effort was structurally incapable of producing success?” Not “what went wrong?” but “what would have had to be true for this to go right, and which of those conditions were missing?”

This is not natural. It requires a disposition most people do not have and most environments do not reward: **the willingness to treat one’s own work history as data rather than narrative.** To look at a decade of forgettable projects and say, not “I was unlucky,” but “I was inside these systems and I can describe exactly what each of them lacked.”

That disposition can be partly innate, but it is mostly trained. Study helps. Reading about production, design, organizational behavior, systems thinking, and market dynamics gives the person a vocabulary for what they are observing. Without that vocabulary, the observations stay intuitions: accurate but inarticulable. With it, the observations become models. Models can be tested, communicated, and applied.

The sequence is: **study gives the vocabulary, experience provides the cases, comparison generates the insight, small experiments test it, and stronger external standards correct the parts you cannot see yourself.** Run that sequence for a decade inside failing organizations and you may end up with a body of practical knowledge about what makes creative work succeed or fail that is very hard to acquire any other way.

6. **Why media and entertainment is special**

There is one feature of media and entertainment (games, film, television, music, publishing) that makes this learning path viable in a way it would not be in aerospace engineering or pharmaceutical research.

The excellent artifact is publicly available.

A person who has never worked at a great game studio can still play great games. They can study them with the same tools they use to build: a game engine, a debugger, a sketchpad, a spreadsheet. They can reverse-engineer the pacing, the economy, the feel, the information hierarchy, the onboarding, the difficulty curve. They can compare it, feature by feature, decision by decision, with the games they helped build that didn’t work.

This is not possible in most industries. An engineer at a mediocre car company cannot disassemble a competitor’s vehicle at the level of design intent. A researcher at a struggling lab cannot observe the daily process of a lab that publishes breakthroughs. The artifact and the process are both hidden.

In games, the artifact is right there. The process is hidden, but the artifact is a dense encoding of the process. A well-designed game is a fossil record of thousands of decisions, and a person with enough experience building games can read that record. They can see where scope was cut, where a mechanic was iterated into something sharp, where a compromise was made and what it cost.

This closes the loop. The person working inside a failing studio has direct experience of what doesn’t work and public access to what does. Combined with study and reflection, those two datasets

produce a living map of the distance between where they are and where the work needs to be.

7. **Where taste, originality, and voice actually come from**

Three capabilities matter more than almost anything else in creative work: taste, originality, and voice. All three share a property that is rarely discussed honestly: **proximity to excellence is not sufficient to develop them, but distance from excellence is not sufficient either.** They emerge from a harsher combination: exposure to strong work, dissatisfaction with weak work, and repeated attempts to close the gap.

Taste is not the ability to recognize good work. Anyone can do that after enough exposure. Taste is the ability to feel, early and accurately, that something is *not yet good enough*, and to know which direction “better” lives in. That feeling is often sharpened by years where nothing quite meets the bar, where the gap between what the work is and what it should be is so wide that the person who cares has to build their own instruments for measuring it. But those instruments still need calibration. Without recurrent contact with truly strong work and people capable of naming the difference, dissatisfaction can harden into private preference rather than taste.

Originality is the refusal to accept inherited solutions when those solutions do not fit the problem. A person trained inside a great studio inherits a toolbox of proven approaches. Those approaches work, which is why they are dangerous: they solve the last studio’s problems, not the next one’s. The person who never had proven approaches to reach for, and had to reason from first principles because nothing around them was worth copying, develops a different relationship with convention. They are not original because they are trying to be. They are original because **imitation was never a sufficient strategy.** But originality without correction drifts easily into idiosyncrasy. What makes it valuable is not merely difference, but useful difference proven against reality.

Voice is the hardest to define and the easiest to recognize. It is the quality in someone’s work that is irreducibly theirs. Institutions

often sand it off. Every review, every template, every house style, every approval chain optimizes for consistency, and consistency can flatten voice. But solitude alone does not produce voice either. Voice emerges when a person is alone with the problem long enough to stop solving it only the way they were taught, and then returns with something strong enough to survive contact with other people.

This is what the successful studio, despite all its advantages, structurally cannot provide. **A disciple who never leaves the master's studio becomes a copy of the master.** They learn the master's taste, the master's solutions, the master's voice. They may execute at an extraordinarily high level. But when they are asked to lead, to set the bar instead of meeting it, they reach for the master's instincts because those are the only instincts they have.

The disciple who leaves and wanders (working in bad studios, shipping forgettable projects, operating under constraints the master's studio would never tolerate) has a different experience. The wandering is uncomfortable. It looks like a downgrade. But it is during the wandering that the disciple stops borrowing and starts building. They discover what they actually care about, as distinct from what they were trained to care about. They develop standards that are *theirs*, not inherited. They build taste by being surrounded by things that offend it, originality by operating where the conventional answers don't work, and voice by being alone with problems long enough that no one else's voice is available.

The people who eventually build the great studios are almost never people who spent their entire careers inside existing great ones. They are people who spent time inside systems that were not good enough, knew why they were dissatisfied, and eventually got the opportunity to build something that resolved it. The dissatisfaction was not a phase. **It was the training.**

Working near excellence often refines execution. Working far from it, while studying it from the outside, can **build the judgment that decides what is worth executing.** One tends to produce reliable contributors. The other can produce people capable of setting direction. Both are necessary. Neither is sufficient alone.

8. Contrasting environments

Not every environment teaches, even to the attentive. Three common patterns suppress learning entirely.

8.1. Where learning is not required

Large, profitable studios with established franchises often operate in a regime where the existing process is good enough. Revenue arrives. Reviews are acceptable. The team knows how to ship. In this environment, individual learning is optional. The machine runs. A person can spend years contributing competently without ever being forced to examine why the machine works or whether it would work under different conditions. The learning that happens is operational (how to navigate the org, how to meet the bar), not structural. The person improves at their function without developing transferable understanding of the system.

8.2. Where learning is institutionalized

Some organizations invest heavily in formal learning: mentorship programs, internal talks, post-mortems with structured templates, reading groups. These are genuinely better than nothing. But institutionalized learning has a ceiling. The institution decides what is worth learning, and that decision is shaped by the institution’s own biases. A studio that has never successfully iterated a core loop will not build a post-mortem process that reveals that gap. The learning infrastructure reflects the organization’s existing model of what matters, which is precisely the model that needs to be questioned.

Institutional learning also tends toward the transmissible: process improvements, pipeline efficiencies, tool upgrades. The harder lessons (taste, scope discipline, the courage to cut, the instinct for when a game is working) resist formalization. They live in judgment, not procedure. Judgment is trained by consequence, not curriculum.

8.3. **Where learning is prevented**

The most damaging environment is one where learning is structurally punished. Studios where questioning the direction is read as disloyalty. Where the founder's vision is not a hypothesis but a decree. Where feedback flows upward only when it is positive, and problems are reframed as execution failures rather than direction failures. In these environments, the attentive person still sees. But they cannot test what they see, because every experiment is insubordination and every question is a threat. The knowledge accumulates but cannot be validated. Over time, the person either leaves or stops looking.

9. **What this means for hiring**

The implication for hiring (https://gerolds.github.io/posts/hiring/) is direct. A track record of failed or mediocre projects is not, by itself, evidence that the person is mediocre. It may be evidence that the person has operated inside mediocre systems. The question is not *what did they ship?* but *what did they learn, and can they articulate it?*

A candidate who can describe, with precision, why each of their past projects failed, not with bitterness or blame but with structural clarity, has demonstrated something rare: the ability to extract signal from noise and to build a model of quality from its absence.

This person may lack polish. They may not have the vocabulary of someone trained inside a prestigious studio. They may not pattern-match to the archetype of a "senior" hire. But they may carry something more valuable: a **generative understanding** of what makes creative work succeed, derived not from imitation but from years of watching it fail and figuring out why.

The person who grew up in excellence and the person who grew up studying its absence are not interchangeable. They have different strengths, different blind spots, different failure modes. But dismissing the second because their résumé lacks prestige is one of the most common and most expensive mistakes a small studio can make.

10. **The synthesis**

Excellence is not a thing you absorb by standing near it. But it is also not a thing you derive from failure alone. It develops by studying the conditions that produce it and the conditions that prevent it, then testing those models against standards harsh enough to expose what you still cannot see.

A career spent inside failing organizations, combined with study, reflection, and the discipline to run small experiments, can produce a practitioner who understands those conditions better than someone who spent the same years inside a successful organization and never had to ask why it worked.

The failure-trained person’s knowledge is hard-won, empirically grounded, and structurally deep. It is also incomplete until it has survived stronger feedback than failure usually offers. And it is invisible on a résumé. It does not come with brand names, shipped titles that sold millions, or the social proof that makes hiring committees comfortable.

That is the claim: **a track record of shitty jobs can produce an exceptionally capable person.** Not despite the shitty jobs. *Partly because* of them. The dysfunction was part of the curriculum. The person who studied it honestly and rigorously, and then sought out better standards than the local culture could provide, learned things no successful studio can teach directly and no failing studio can teach completely.

The only question, for the person and for anyone evaluating them, is whether they actually did the studying, and whether they found ways to test what they learned against reality. The environment provided the material. Whether it became an education depends on whether someone was paying attention and whether they kept climbing toward feedback that could still prove them wrong.

11. Appendix: further reading

11.4. On learning from failure and negative space

Karl Weick, *The Social Psychology of Organizing* (2nd ed., 1979). Weick's concept of "sensemaking" describes how people construct explanations from ambiguous data. The person who learns inside a failing organization is doing exactly this: making sense of dysfunction in real time and building models from noise.

Charles Perrow, *Normal Accidents: Living with High-Risk Technologies* (1984). Perrow argues that in tightly coupled, complex systems, failure is not an aberration but a structural inevitability. The same logic applies to game production.

Amy Edmondson, *The Fearless Organization* (2018). Psychological safety as the precondition for learning from failure. Edmondson's research helps explain why most people inside failing organizations do not learn: the environment punishes honesty.

Matthew Syed, *Black Box Thinking* (2015). A comparison of industries that learn from failure (aviation) and those that don't (medicine). The core argument, that open investigation of failure is the fastest path to improvement, applies directly here.

11.5. On taste, originality, and creative voice

Ira Glass, "The Gap" (interview, 2009). Glass describes the period where a creative person's taste outstrips their ability, where they can see that their work isn't good enough but cannot yet close the distance.

Harold Bloom, *The Anxiety of Influence* (1973). Bloom's theory of poetic development: every strong artist begins as a disciple and must "misread" their predecessors in order to create original work.

Robert Henri, *The Art Spirit* (1923). Henri's teaching philosophy: the student must eventually reject the teacher's solutions and find their own, not through rebellion but through honest encounter with the material.

Brian Eno, "The Big Here and Long Now" (1996, essay) and various interviews on generative processes. Eno's distinction

between “genius” as individual talent and “scenius” as the intelligence of a creative community is useful here.

11.6. **On expertise outside prestigious institutions**

K. Anders Ericsson, *Peak: Secrets from the New Science of Expertise* (2016). Ericsson’s concept of “deliberate practice” requires feedback, reflection, and progressive difficulty. Failing organizations may provide discomfort and many opportunities for reflection, but they often do not provide the kind of precise, high-quality feedback that excellence requires. The attentive person has to seek or construct that missing feedback elsewhere.

Nassim Nicholas Taleb, *Antifragile* (2012). Systems that gain from disorder. Useful for thinking about the difference between someone strengthened by exposure to dysfunction and someone trained under unusually stable conditions.

James C. Scott, *Seeing Like a State* (1998). Scott’s concept of “mētis,” practical knowledge that resists formalization and is acquired only through direct engagement with a specific environment.

Matthew B. Crawford, *Shop Class as Soulcraft* (2009). Crawford argues that manual trades develop a form of intelligence that abstract knowledge work does not: the constant confrontation with material that resists your intentions.

11.7. **On institutional momentum and the limits of learning inside success**

Clayton Christensen, *The Innovator’s Dilemma* (1997). Successful organizations develop processes and values optimized for their existing business. Those processes become filters that reject disruptive possibilities.

Jerry Harvey, “The Abilene Paradox” (1974). Groups collectively pursue a course of action that no individual actually believes in, because each assumes the others are committed.

Further reading within this textbook:

- *Hiring for Games* *(https://gerolds.github.io/posts/hiring/)* — how to evaluate candidates whose résumés lack prestige but whose understanding runs deep
- *Studio DNA* *(https://gerolds.github.io/posts/studio-dna/)* — reading what a team is adapted for, and recognizing where institutional reflexes substitute for transferable skill
- *The Dream That Won't Compound* *(https://gerolds.github.io/posts/the-dream-that-wont-compound/)* — the studio-level pattern of surviving without improving, and why the funding structure prevents the learning this essay describes
- *The Price of Seeing Clearly* *(https://gerolds.github.io/posts/the-price-of-seeing-clearly/)* — how the mindset shift actually happens: not by decision, but by the collapse of the old model
- *The Studio Primitive* *(https://gerolds.github.io/posts/studio-os/)* — the Promise → Proof → Cut → Ship loop as a structure that converts individual insight into repeatable process
- *Prototyping the Loop* *(https://gerolds.github.io/posts/prototyping-the-loop/)* — finding the invariant: the small-experiment methodology this essay argues is uniquely available in failing organizations
- *Wicked Problems* *(https://gerolds.github.io/posts/wicked-problems/)* — the failure patterns that become visible only to someone who has lived inside them long enough to name them
- *Why Things Compound* *(https://gerolds.github.io/posts/why-things-compound/)* — the compounding mechanism that the failure-trained person understands by its absence
- *Making the Thing* *(https://gerolds.github.io/posts/making-the-thing/)* — deriving a coherent game from a governing commitment, the skill the wandering practitioner develops by watching incoherence repeatedly

---

Drafting assistance: Claude. All claims mine; errors my responsibility.