Failing Confidently

JOSH SOLOMON · KNOWN QUANTITY LLC · NOVEMBER 2026

> Other people's forecasts of you are usually accurate about your past and wrong about your ceiling, and failure gives them more evidence every time. An evidence-based case for betting on your own estimate of what is possible, expecting to lose often, paying for it yourself and refusing to let the losses lower the bet, with an honest account of where overconfidence and survivorship bias cut against it.

Other people's ceiling

For most of my life the people responsible for me believed in me. That is the strange part. Teachers, administrators, managers, the adults in charge. They believed in what they thought was possible for me. They never believed in what I thought was possible, and the gap between those two numbers crushed me year after year.

The record gave them reasons. I was nearly expelled in middle school, nearly kicked out of summer camp, nearly kicked out of college more than once. After college I quit a job, lost money trying to trade for a living, started a company that failed, and came back to corporate work that did not want me. Every one was a data point that lowered someone's estimate of me. The estimates were reasonable, which is what made them hard to argue with.

What I have learned since, mostly by being wrong, is that the failures were not the problem. The problem was letting each failure, and each person's reasonable forecast, set the size of the next bet. The years that went well were the ones where I bet on my own number, lost, and placed the same bet again with a better method.

That is the argument of this essay. Failing is not a virtue and it is not a strategy. But the people who end up somewhere unusual fail about as often as everyone else; what separates them is what they do with the failure and whose ceiling they keep. Fail confidently: commit to your own estimate of what is possible, expect to be wrong often, pay for it yourself, and do not let the loss or anyone else's lower ceiling revise the estimate down.

The premise

By failing confidently I mean four things together, and the definition only works with all four.

First, you commit to your own estimate of what is possible, not the one other people hold for you. Their estimate is usually built from your record. Yours includes things the record cannot see: what you intend to do differently and how fast you can learn.

Second, you expect to be wrong often. Not occasionally. A bet on the edge of what you can do should fail most of the time, and if it does not, the estimate was too low.

Third, you pay for it yourself. Your time, your money, your comfort. Not your employer's budget, not your family's savings without their consent, not your investors' money on a story you do not believe.

Fourth, the failure revises the method and never the ceiling. You lost; something about how you tried was wrong; you change that. What you do not do is conclude that the people with the lower number were right about the number.

It is not recklessness. Recklessness is betting what is not yours, or betting without expecting to lose. It is not denial either. Denial refuses to look at the failure, and failing confidently depends on looking at it more closely than anyone else will. And it is not "fail fast," the startup slogan that treats failure as the goal. Failure is the cost. The goal is the thing on the other side of it.

The last essay argued that caring about the verdict makes you worse at the work, and that the fix is to keep the effort and drop your stake in the result. This is the same move pointed at a different verdict. There, the verdict was whether a post or a design worked. Here, it is the verdict that other people and your own losses pass on you, and the question is whether you let it set your ceiling.

What the evidence says

The largest study of failure I know of separates winners from losers on something other than how much they failed. Yian Yin, Yang Wang, James Evans and Dashun Wang followed repeated attempts in three domains: 776,721 NIH grant applications from 139,091 scientists, 58,111 startups, and 170,350 attacks by 3,178 terrorist groups. In all three, the people who eventually succeeded did not try more or less often than the people who never did, and they started at statistically indistinguishable levels. What differed was the dynamics. For the eventual winners, the time between attempts shrank with each failure, following a power law, and the two groups could be told apart as early as the second failure. Their model ties progress to how much of each past attempt you carry into the next (Yin et al., Nature 2019; full text). Failing a lot is not the signal. Failing faster, while keeping what the last attempt taught you, is.

The second finding is the closest thing to a clean experiment on a setback. Yang Wang, Benjamin Jones and Dashun Wang compared early-career scientists whose first major NIH grant applications landed just below the funding line with those just above it: 623 near misses against 561 narrow wins, essentially identical on paper. The near miss cost something real. It raised the chance that a scientist left the NIH system over the next decade by about 12.6 percent. But the near misses who stayed went on to outperform the winners: 16.1 percent of their papers in the next five years were hits, in the top 5 percent of citations, against 13.3 percent for the narrow winners. The gap held when the authors removed the weakest 12.6 percent of winners to mimic the attrition (Wang, Jones and Wang, Nature Communications 2019; Kellogg Insight). The same setback that pushed some people out made the ones who kept their bet better.

Third, most of us do not learn from our own failures, and the reason is ego. Lauren Eskreis-Winkler and Ayelet Fishbach ran five experiments with more than 1,600 people on simple two-choice questions, where telling someone they were wrong tells them exactly what is right. People still learned less from failure feedback than from success feedback, even when learning was paid and the task was made easier, and they remembered less of what they had chosen. When the same people watched someone else fail, they learned as much as from someone else's success (Eskreis-Winkler and Fishbach, Psychological Science 2019, via UChicago). Failure contains the information. The self blocks it.

Fourth, failure by itself is not a credential. Paul Gompers, Anna Kovner, Josh Lerner and David Scharfstein tracked 8,753 venture-backed founders in 3,796 companies from 1986 to 2000, counting success as going public or filing to. First-time founders succeeded 21 percent of the time. Founders whose previous venture had failed succeeded 22 percent of the time, no better. Founders who had already succeeded hit 30 percent (Gompers et al., Journal of Financial Economics 2010). Prior experience of any kind does help in plainer businesses: across every taxable establishment opened in Texas from 1990 to 2011, owners who had started a business before kept the next one open longer, even comparing the same person across businesses, which the authors read as learning by doing (Lafontaine and Shaw, NBER 2014). The honest summary is that failing is neutral. What you carry out of it is the variable.

Fifth, other people's expectations of you matter less than the famous story says, and they are mostly accurate. The Pygmalion study, where teachers told that random students were about to bloom saw those students gain IQ points, is weaker than its reputation; as Lee Jussim and Kent Harber's review lays out, the average gain was about 4 points, the large effects were confined to the first two grades, and a handful of implausible scores drove much of it. Across naturalistic studies, the self-fulfilling part of teacher expectations runs around r = .1 to .2, does not accumulate much over time, and is dwarfed by accuracy: roughly three quarters of what expectations predict about test scores is the teacher being right. The one place the effect may be large is among low achievers and stigmatized groups, where several studies found it two to several times stronger, though the authors say plainly that this rests on few studies (Jussim and Harber, Personality and Social Psychology Review 2005). So people are not usually wrong about you. They are forecasting your average from your record, and they are usually close. The damage, if there is any, lands hardest on the people whose record is worst.

Sixth, one person can be enough. Emmy Werner followed every one of the 698 children born on the Hawaiian island of Kauai in 1955 to age 32. About a third, 201 children, were high risk: poverty, birth complications, troubled homes. Two thirds of those had serious problems by adolescence or early adulthood. One in three, 72 of them, grew into what Werner called competent, confident and caring adults, and they shared a feature: "at least one person in their lives who accepted them unconditionally," often a grandparent, an older sibling, a teacher or a mentor. Many of the troubled ones recovered later, through community college, the military or a church, getting what Werner called a second chance at self-efficacy in adulthood (Werner, Development and Psychopathology 1993). It is one cohort on one island and it is correlational. It is also three decades of data on everyone born there that year, and the variable it points to is not a program. It is a person.

Seventh, believing you can do the specific thing helps, and believing in yourself in general helps much less. Alexander Stajkovic and Fred Luthans pooled 114 studies and found a weighted correlation of .38 between task-specific self-efficacy and work performance (Stajkovic and Luthans, Psychological Bulletin 1998). The evidence for growth mindset is much weaker than its popularity. Across 273 effect sizes and 365,915 students, mindset barely related to achievement, and across 43 intervention effects with 57,155 students the overall benefit was weak, with somewhat larger gains for low-income students (Sisk et al., Psychological Science 2018). The best-run test, a national experiment with more than 12,000 ninth graders and two 25-minute online sessions, raised lower-achieving students' core grades by 0.10 grade points and cut D and F averages by more than 5 percentage points (Yeager et al., Nature 2019, via EurekAlert). Real, cheap and small. Confidence in your ability to do a particular hard thing tracks performance. A poster about believing in yourself does not.

Eighth, and this is the one that cuts against me: overconfidence is real, persistent and expensive. Don Moore and Paul Healy separated it into three kinds and found that people overestimate their own performance on hard tasks, and that overprecision, being too sure of your own estimate, is the most persistent of the three (Moore and Healy, Psychological Review 2008). In the 1988 founder survey cited in the last essay, a third of new business owners rated their own odds at 10 out of 10 (Cooper, Woo and Dunkelberg, 1988). Among independent inventors told by an expert program that their project should stop, 29 percent kept spending money and 51 percent kept spending time, and the optimists spent 166 percent more than the pessimists after the advice (Åstebro, Jeffrey and Adomdza, Journal of Behavioral Decision Making 2007). And the self-employed in the United States earn less: a median gap of 35 percent after ten years in business, which Barton Hamilton attributed to people accepting lower pay for independence (Hamilton, Journal of Political Economy 2000). Any honest case for betting on your own number has to survive these.

Case study: a record of other people's forecasts

This is one person's record, so treat it as a case, not a proof. It is useful because it contains the same pattern many times over: a reasonable forecast from the people in charge, a failure that confirmed it, and occasionally someone who held a different number.

The early entries are about rules. I was nearly expelled from middle school for bad language with a teacher. The school was small, and afterward I was not allowed at classmates' houses. One friend and that friend's family stood by me anyway. I was nearly kicked out of summer camp. In college I was nearly kicked out several times for breaking rules. None of this is unusual, and every adult who made a forecast from it was forecasting sensibly. Jussim's review would predict they were mostly right about where a kid like that was headed.

One entry does not fit. I wrote an exam on the wrong subject, because I was so excited. A teacher who believed in me gave me a second chance, and I aced it. That is the Wang, Jones and Wang result at the scale of one person: a near miss, a second attempt, a better outcome. The difference was one person who

held a higher number than the record supported.

The cost of the record was concrete. I could not study abroad because of my standing with the school. I wanted to go to Thailand. I still have not been.

The next entries are about work. My first job out of college was far too hard for me, and I made it through. Then I quit, tried to trade for a living during covid, lost money and got depressed. I took a job that was unfulfilling and a step down. I started a company that did well at first, quit the job for it, and then watched it fail worse than anything before it, while I was traveling in Mexico and Germany. Coming back to corporate work after that was nearly impossible. The market was not being unfair. Gompers would predict exactly that: a failed founder is, on average, no better a bet than someone who never tried.

Two of those failures fit the overconfidence literature, and I want to say so plainly. The trading looks like Moore and Healy's overprecision: being too sure of my own read. The company was the founder's blind spot from the last essay. I paid for both myself, which is the part of the definition I got right. The rest of the definition is what I had to learn after.

Then one company gave me a chance. I knew basically nothing when I started and walked away knowing everything I now use. That is Lafontaine and Shaw's learning by doing, and Werner's second chance in adulthood, in one entry. Then I was laid off, which turned out to be the best thing that happened to my work, because for the first time nobody's forecast stood between my estimate and the attempt. In the last month I have shipped a game studio with dozens of browser games, a consulting practice and these essays, and run the channels described in the last essay.

There were personal losses along the way that made all of it harder. And there is my wife, the only person who has fully believed in my version of what is possible, not the version my record supports. She also challenges me, which matters as much. I am happy now.

What the case does not show: that confidence caused any of it. It is a record written by someone for whom things turned out well, which is the definition of survivorship bias. What it shows is narrower. The forecasts were accurate about my record and wrong about my ceiling, and the people who changed the outcome were the few who held a higher number than the record justified.

Why other people's ceilings stick

Put the findings together and the mechanism is not mainly about other people.

Their forecast is built from your record, and it is usually right about your average. That is what Jussim found. It is also why the forecast cannot see the tail. An estimate built from base rates puts a person with a bad record near the bottom and is correct most of the time. The few it is wrong about are, by construction, the ones it had no way to see.

The forecast hurts because it hands you an exit. A near miss in the Wang data pushed some scientists out of the system entirely. Nobody forced them. When the people around you hold a lower number and you fail, their number becomes the easiest explanation available, and accepting it ends the pain immediately. It also ends the bet.

Failure then blocks the very learning that would justify a higher number. This is Eskreis-Winkler and Fishbach. Your own failure is an ego threat, so you look away from it, and the information that would make the next attempt better never arrives. Yin's winners are the people who did not look away: each attempt reused more of the last, so the attempts got faster and better. That only happens if the failure is something you can examine without it saying anything about who you are. The last essay was about exactly that condition.

And overconfidence sits on the other side of the road. Moore and Healy's people overestimate themselves on hard tasks and are too sure of their estimates. So the problem is not simply "believe in yourself more." A person who raises his estimate of every attempt is the inventor who keeps spending after the expert says stop.

The way through is to separate two numbers that usually get collapsed into one. One is the ceiling: what is possible for you over many attempts, if each one teaches you something. The other is the odds of this attempt. Other people's forecasts are mostly about the odds, and on the odds they are usually right. Failing confidently means accepting their odds, often worse than their odds, and keeping your own ceiling. Overconfidence is the error of being too sure this attempt will work. Failing confidently expects this attempt to fail and makes it anyway. Those are different bets, and the evidence penalizes the first one.

Three old versions of the same idea

None of this is new. Three traditions got there first, and each one is more careful than the version on the poster.

Karl Popper built a philosophy of science on failing confidently on purpose. A scientist should make the riskiest claim he can, one that might easily turn out false, and then try hard to break it: "Every genuine test of a theory is an attempt to falsify it" (Popper, Conjectures and Refutations, 1963, p. 36). A theory that cannot fail says less. The detail people forget is the one that matters here. Popper's scientist, according to the Stanford Encyclopedia's account, does not abandon the present theory "until such time as he has a better one to substitute for it" (Thornton, Stanford Encyclopedia of Philosophy). A refutation revises the theory. It does not revise the ambition to explain the thing. That is the fourth part of the definition: the failure changes the method, not the ceiling.

Proverbs puts it in one line: "the righteous falls seven times and rises again" (Proverbs 24:16, ESV). It is usually quoted as a line about moral failure. The commentaries disagree. Ellicott reads it as falling into trouble, "not sin, as is often supposed," and the Pulpit Commentary notes that the Hebrew verb is not used for moral falling and that "seven times" just means often. The verse before it warns a wicked man lying in

wait near a good man's home. So the proverb is not about guilt. It is about a person who meets calamity often, in front of people who expect him to stay down, and gets up anyway. It does not promise he will not fall again.

Samuel Beckett gave the poster its favorite line: "Ever tried. Ever failed. No matter. Try again. Fail again. Fail better." It comes from Worstward Ho, a late, fragmentary prose piece from 1983, and the next lines are "Fail worse again. Still worse again." Slate's account of how the phrase became a Silicon Valley mantra calls the full text "worse than useless as a motivational aid" (Slate, 2014). The bleakness is the useful part. Beckett promises nothing on the other side. "No matter" is the point: the failure is not allowed to mean anything about whether you try again. That is not optimism. It is refusing to let the result set the terms.

What the three share is the thing this essay is about. Make the strong claim. Expect it to fail, often. Do not let the failure, or the people waiting for it, decide whether you make the next one. None of them says failing is the goal.

How to fail confidently without fooling yourself

The evidence does not say fail more. It says carry more out of each failure, keep your own ceiling, and stay honest about your odds. Each step below is tied to the research that says why.

Write down two numbers before you start. The ceiling: what you think is possible over many attempts. The odds: your honest chance on this one. Make the odds lower than feels comfortable. Moore and Healy say you will overestimate on hard tasks; correct for it in writing. Then attempt it anyway. If someone tells you the odds are bad, you can agree with them and lose nothing.

Shorten the time between attempts, and reuse the parts that worked. This is the Yin result directly. The winners were not the ones who failed most. They were the ones whose gaps between attempts shrank because each try carried forward more of the last. After a failure, list what you will keep, not just what went wrong, and start the next attempt before the list goes cold.

Review your failures as if they were someone else's. Eskreis-Winkler and Fishbach's subjects learned from other people's failures as well as from successes. The third-person review from the last essay is the same tool: write the post-mortem about "him," and the information gets through.

Be kind to yourself about the failure, not proud of yourself in general. In Juliana Breines and Serena Chen's experiments, people who failed a hard vocabulary test and were then prompted toward self-compassion studied about a third longer before a retest than people prompted toward self-esteem and about half again as long as people given no prompt, and across four experiments it raised the motivation to improve (Breines and Chen, Personality and Social Psychology Bulletin 2012). The samples were small, around 70 to 90 per study. The direction is the useful part: forgiving the failure keeps you looking at it.

Attempt before you are ready. In learning research this is called productive failure. A meta-analysis of 53 studies and 166 comparisons found that letting students struggle with a problem before teaching them the method beat teaching first, with an effect of g = 0.36, and up to 0.58 when the struggle was designed well (Sinha and Kapur, Review of Educational Research 2021). It needs some prior knowledge to work. The extension to adult work is mine, not theirs: the failed first attempt is part of how you learn the method, not proof you should not have tried.

Pay for it yourself, and cap what you can lose. Hamilton's 35 percent and the inventors who kept spending are the price of getting this wrong. Decide in advance what the attempt is allowed to cost, in money and in months, and hold to it. Failing confidently with someone else's rent is not confidence.

Find one person who holds your number, and make sure they will also argue with you. Werner's resilient children had someone who accepted them unconditionally. That is necessary and not sufficient. The person who only agrees with you is the inventor's optimism with a second voice. You want someone who believes in your ceiling and will tell you when your odds are bad.

The objections I take seriously

"This is survivorship bias. You are writing it because things worked out." Yes, and it is the strongest objection. I am a case of one, and the people who placed the same bets and did not recover are not writing essays. That is why the argument rests on studies that include the people who quit: Yin's non-winners, Wang's near misses who left, Gompers' founders who failed before. In those data the failures are counted, and the finding still holds that what separates the outcomes is how people use a failure, not whether they had one.

"Overconfidence ruins people. You are telling them to ignore the experts." I am telling them to agree with the experts about the odds and disagree, if they have reason to, about the ceiling. The inventors who kept spending after being told to stop are the warning, and the definition is built to exclude them: they were overconfident about the attempt itself, and they kept paying after the cost should have been capped. If an expert says your project is bad, the confident response is usually to stop that project and keep the ambition, which is what Popper's scientist does with a refuted theory.

"Other people's expectations were probably right about you." About my average, almost certainly. That is what Jussim found, and a forecast built from my record would have been accurate for most people with that record. The argument is not that forecasters are wrong. It is that a forecast built from the past is the wrong number to set your ceiling by, because it has no way to include what you will do differently, and that the cost of listening to it falls hardest on the people with the worst records.

"Most people do not learn from failure, so this just means more pain." That is the Eskreis-Winkler finding and the Gompers 22 percent, and it is correct. Fail confidently without changing the method and you get the pain without the benefit. That is why the practical section is mostly about how to look at a failure, not how to collect more of them.

"Easy to say with a safety net." True, and I have had one: my wife's belief is a safety net of its own. The definition includes paying for the bet yourself precisely because the bets that ruin people are the ones paid for by someone who did not agree to them. The honest version is: bet the amount you can lose, and if that amount is small, make smaller bets more often. Yin's result is about the pace of attempts, not their size.

"Mindset research is a fad with tiny effects." Mostly right, which is why I have not leaned on it. What holds up is narrower: task-specific confidence tracks performance, and how you treat a failure changes whether you learn from it.

What to do on Monday

1. Pick the thing you have stopped attempting because someone's forecast, or your record, said it was out of reach. Write two numbers on paper: the ceiling you actually believe in, and your honest odds on the next attempt. Make the odds low.
2. Decide what the next attempt is allowed to cost, in money and in weeks. Write it under the two numbers and date it.
3. Go back to your last failure and write a short post-mortem about it in the third person. Two lists: what he got wrong, and what he should keep. The second list matters more.
4. Start the next attempt sooner than you did the last one, and build it from the keep list. The goal is a shorter gap than last time.
5. Tell one person your ceiling, not your odds. Pick someone who will believe it and argue with you about how you are going about it. If you do not have that person, look for them before you look for anything else.

You will probably fail. Fail on your own number, and make the next attempt anyway.

Sources and a note on method

Every figure is linked where it appears, to the primary study or the most direct report available. Where a claim rests on one study or one cohort I say so; where it rests on a meta-analysis I give the number of studies. A few numbers come from a direct report rather than the paper, because the paper was not open to me (Eskreis-Winkler and Fishbach, Yeager, and the Wang, Jones and Wang hit-paper rates); those link to the report I checked.

A few things that are commonly said were left out on purpose because the evidence does not carry them: the original Pygmalion effect as a dramatic IQ boost (about 4 points, two grades, a handful of implausible scores); the claim that growth mindset broadly raises achievement (the meta-analytic effects are small); the 10,000-hour rule and strong claims about deliberate practice explaining most of expertise; the statistic that most successful founders failed first (in venture data, a prior failure predicts nothing); and the quote

"Fall seven times, stand up eight" attributed to Samurai or Zen masters, which I could not trace to a primary source and replaced with a verse that could be checked. My own record is a case of one and is presented as such.

The research brief behind this essay, with a strength-of-evidence rating on every item and the list of claims cut, is available on request. The first two essays in this series are The Ego Death of Society and The Utility of Not Caring.

I used AI to research and draft this, under my direction, with a standing instruction to tell me where I was wrong before telling me where I was right, and every claim was checked against its source before it went in. Given the subject, being told I was wrong early and often was the point.