The Utility of Not Caring

JOSH SOLOMON · KNOWN QUANTITY LLC · NOVEMBER 2026

> What the tax costs, where it comes from, and how to stop paying it without becoming someone who does not care about anything: one founder's record of running the channels he hated better than the game he loves, the research on why, and three old traditions that got there first.

I hate all of it, and it worked anyway

I do not like LinkedIn. I do not like TikTok. I do not like the voice people use on either one, the fake vulnerability, the lessons learned, the hustle. I have never once wanted to be an influencer. In October I needed both of them to work, because I had a company to launch and no budget, so I did the only thing I knew how to do with something I did not care about: I treated it as data.

No opinions. What does the platform reward? Hook in the first second, a question at the end, post at these times, this length, these tags. Measure, change one thing, measure again. Within a few weeks the game studio's clips passed a million combined views on TikTok, with one clip at 332,000, and a single LinkedIn post about the company reached 188,000 people. Not luck. I had no feelings about any of it, so I did exactly what the numbers said, and the numbers were right.

Then I looked at the things I love. The game I have been building for a month, where I argue with every design decision and ship late because it is never good enough. The guitar channel I started this week and spent an hour on the avatar for. The essays. Everything I care about, I run worse. I second-guess, I protect, I take the data personally, I keep the version I like over the version that works.

That is the whole observation, and I think it is bigger than me. Caring is not free. It is a tax on judgment, and the tax is highest exactly where we most want to be right. This essay is about what the tax costs, where it comes from, and how to stop paying it without becoming someone who does not care about anything.

The premise

By not caring I do not mean apathy, and I do not mean cynicism. I mean something narrower: having no stake in which answer turns out to be true. When I posted the first clip I did not care whether a 0.8-second hook beat a 2-second hook. I ran both. Whichever won, won. I had no pride in the loser, so I could drop it the same afternoon.

Compare that with the thing you love. There, you have a stake in the answer before the test runs. You want the version you built on a Saturday to be the one that works. You want the song you wrote to be the one people like. You want the plan you defended in the meeting to be the one that ships. So you read the data looking for the result you want, you run the test one more time when it goes against you, and you give the thing you love another chance that the thing you did not care about would never have gotten.

That is the tax. It shows up in four places: you see less (you filter the evidence), you decide slower (every change feels like a loss), you quit later (the sunk cost is also your identity), and you perform worse under pressure (because now it matters). Each of those has a research literature behind it, and the next section is a tour of it.

The argument is not that passion is bad. Passion is why you show up. The argument is that passion should pick the project and then get out of the room while the decisions are made. The founder who built a million-view channel out of a platform he despises did not have better taste or better tools than the founder who stalls on the game he adores. He had less skin in the answer. The rest of this essay is about how to manufacture that condition on purpose, for the things that matter.

What the evidence says

The cleanest demonstration has my name on it, which I did not plan. Psychologists call it Solomon's paradox, after the king who judged everyone's disputes wisely and ran his own life into the ground. In three experiments with 693 people, Igor Grossmann and Ethan Kross found that we reason more wisely about a friend's problem than about an identical problem of our own: we see the limits of what we know, we weigh the other side, we look for compromise. Then they asked people to think about their own problem from a distance, as if watching someone else, and the gap closed (Grossmann and Kross, Psychological Science 2014). In a preregistered follow-up with 555 adults, a month of writing a daily diary in the third person raised the quality of people's reasoning about their own conflicts (Grossmann et al., 2021). The posts I did not care about were, in effect, someone else's posts. That is why I could run them well.

The second finding is about what caring does to money. Everyone knows the sunk-cost effect. Fewer people know what actually drives it. A meta-analysis of 166 samples on escalation of commitment, the habit of pouring more into a failing course of action, ranked the causes. Sunk costs themselves came in at a correlation of .24. Ego threat came in at .47, twice as strong. Having personally made the original decision roughly doubled the escalation in the classic experiment, where students who had chosen a failing division then gave it $13 million of a $20 million follow-up budget, against about $9.5 million from students who inherited the choice (Staw, 1976; Sleesman et al., Academy of Management Journal 2012). It is not the money we cannot let go of. It is the fact that it was ours.

Third, the weighting itself bends under stakes. Loss aversion, the finding that losses hurt about twice as much as equivalent gains, is real, but its size moves: about 2 for $100-scale gambles, only 1.12 for gambles of a few euros, and it shrinks when people know the domain cold (Mrkva et al., 2020, N = 17,720). The more something matters to you, the more your scale is bent, and expertise is the only thing that straightens it.

Fourth, pressure does not sharpen experts, it blunts them. Experienced golfers putt worse when made to attend to their own stroke; novices do not. People with the most working memory lose the most under pressure, because pressure eats exactly the resource they rely on (Beilock and Carr, 2001, 2005). You have probably heard this as an inverted U with an optimal level of stress. That curve rests on twelve mice in 1908 and one of three experiments; the replications failed. The defensible version is simpler and worse: when it matters, skilled people get worse at the skill.

Fifth, and this one surprised me, how much you care barely predicts how well you do. Robert Vallerand's group distinguishes harmonious passion, which you control, from obsessive passion, which controls you and is wired to your self-worth. Across 94 studies, obsessive passion correlated .25 with negative affect and .09 with objective performance, which is nothing. Harmonious passion predicted life satisfaction (.39) and protected against burnout (-.53) and still only reached .10 with performance (Curran et al., 2015). Grit, the popular version of the same idea, correlates .18 with performance across 66,807 people, and the perseverance half does the work while the passion half does almost none (Credé, Tynan and Harms, 2017). Passion is good for you. It is not good for the work.

Sixth, the people most sure about their own thing are the ones building it. Of 2,994 new business owners, 81% rated their own odds of success at 7 out of 10 or better, and a third said 10 out of 10, certain. Asked about businesses just like theirs, only 39% were that confident (Cooper, Woo and Dunkelberg, 1988). Roughly half of new US businesses are gone in five years. The founder's own venture is the one thing the founder cannot see straight.

Seventh, when you stop caring which answer wins and just test, most of your ideas lose. At Microsoft, about one third of the ideas run on its experimentation platform improved the metric they were built to improve; a third did nothing; a third hurt. In mature products like Bing the win rate is 10 to 20%, and Booking.com has reported about one in ten (Kohavi and Longbotham, 2023; Thomke, HBR 2020). These are the best product teams in the world, and they are wrong about two thirds of their own ideas. The only reason they win anyway is that they stopped needing to be right before the test.

And the mechanism underneath all of it is old. Ziva Kunda showed in 1990 that we reason toward the conclusion we want and stop when we have a story we can defend (Kunda, Psychological Bulletin 1990); Slovic's group showed that liking something makes us rate it both safer and more beneficial, more so under time pressure (Finucane et al., 2000); and Stanovich and West showed that intelligence

does not protect you from any of it, with higher numeracy making people more polarized about data on a charged topic, not less (Stanovich and West, 2008; Kahan et al., 2017). Being smart does not fix caring. If anything, it gives you better arguments for the answer you already wanted.

Case study: the things I did not care about, and the things I did

This is one person's record, so treat it as a case, not a proof. It is useful because the same person, with the same tools and the same month, ran two kinds of project side by side.

The things I did not care about were the channels. I had never posted on TikTok. I had never written a LinkedIn post. I find both platforms tiresome and I think most of what is on them is theater. So I ran them like an analyst runs anything he has no opinion about. What does the platform say it rewards? TikTok has said completion matters more than nearly any other signal it discloses, and that follower count is not a direct factor (TikTok, 2020). So: clips short enough to finish, the payoff in the first second, nothing that needs context. Post, read the retention curve, change one thing, post again. Within weeks, the studio's clips passed about a million combined views, the top one at 332,000, and a single post about the company reached 188,000 people on a platform where my profile had been a blank for seven years. Not one of those posts was something I was proud of while making it. Every one of them was a test I was indifferent to the result of.

The things I cared about went differently. The game I have been building since October, the one I think could be the best thing the studio ever ships, is a month in and I still argue with it daily. The first fighting game the AI built in two hours looked bad and I sent it back; that was good judgment. But I also kept art I liked over art that tested better, delayed launches to fix things nobody but me would notice, and read the play data for reasons it was fine rather than for what was wrong. The guitar channel I started this weekend had a logo, a bio and a content calendar before it had a single video. This essay series, which I want to be taken seriously, took longer to start than any channel I treated as a spreadsheet.

The numbers from the trading engine made the same point from a different angle. When the AI went through seven years of my trades, the losses did not come from bad charts. Two tickers I believed in, the two I had the strongest opinions about, accounted for about two thirds of everything I lost; everything else was close to flat. The other big leak was holding options to expiry instead of selling them for whatever they were worth: the expirations lost more than three times what every option I sold before expiry made. Both are the same mistake: caring about a position more than the position deserved. The rules-based engine it built beat my record for one reason: the engine does not care which trade wins. It also did not beat an index fund, which is the honest footnote, and I kept it in.

I want to be careful about what this shows. It does not show that indifference causes virality; there is no peer-reviewed study of creators who treat content as experiments against creators who treat it as expression, and the vendor statistics that circulate about hooks and retention have no traceable source. What the case shows is narrower and, I think, more useful. The same person made better decisions about the thing he did not care about than about the things he did, and the difference was not information, tools or effort. It was stake.

Why the things we love are the hardest to run

Put the findings together and the mechanism is not mysterious. When you care about an outcome, four things happen at once, and each one is a decision-quality problem dressed up as a feeling.

You lose the outside view. The judgment you would give a friend in five seconds is unavailable to you about your own project, because you are inside it. Grossmann's subjects were not dumber about themselves; they were closer. Distance is what wisdom runs on, and caring removes it.

You stop testing and start defending. Once the thing is yours, every piece of evidence gets sorted into for and against, and against gets a harder look. This is Kunda's motivated reasoning doing its ordinary work. The Microsoft number is the antidote and the indictment at once: the best product people in the world are wrong about two of every three ideas, and they only find out because they built a machine that does not care which idea was theirs.

Your scale bends. The loss you fear is weighted heavier than the gain you want, and more heavily the more it matters. So you hold the losing trade, keep the slow feature, protect the design you like. The sunk-cost literature's quiet finding is the one to remember: ego threat drives escalation twice as hard as the money does.

And your hands get worse. Under pressure, skilled people attend to their mechanics and their working memory fills with the stakes. The guitarist who plays beautifully alone and stiffens on camera is not less skilled on camera. He cares more on camera.

Notice what is not on this list. Effort is not the problem. Knowledge is not the problem. Intelligence is not the problem; Stanovich showed it does not protect you, and Kahan showed numeracy can make it worse. The problem is a single variable, your stake in the answer, and it degrades perception, testing, weighting and execution simultaneously. That is why the passion research comes out the way it does: caring predicts how you feel about the work, and almost nothing about how well it goes.

There is a cruel symmetry here. The thing you care about most is the thing you most want to run well, and it is exactly the thing you are structurally worst at running. The channel I despised got my best judgment because it got none of my hope. The game I love gets all of my hope and therefore a worse

version of me. If that is true for a one-person company, it is true for a founder's first startup, a parent's first child, a scientist's own theory and a nation's idea of itself.

Three old versions of the same idea

None of this is new. Three traditions on three continents arrived at it, and each one was careful to say what it is not.

The Stoics drew a line between what is up to us and what is not. Epictetus opens the Handbook with it: within our power are opinion, motivation, desire and aversion; not within our power are the body, property, reputation, office (Epictetus, Enchiridion 1). Health, wealth and reputation they called indifferents, but preferred ones: worth pursuing, never worth staking your good on. The Stoic acts with a reserve clause, "I will sail, unless something prevents it," which Pierre Hadot reads as the key to the whole discipline: you act with all your energy toward the outcome, and you hold the outcome itself loosely (Hadot, The Inner Citadel, ch. 7). That is not apathy. It is full effort with the ego subtracted from the result.

The Gita says it in one line that Indian schoolchildren still memorize. "You have the right to work, but never to the fruit of work." The next clause is the one people forget: "nor should you long for inaction" (Bhagavad Gita 2.47, Easwaran translation). Krishna is talking to a warrior who wants to put down his bow because he cares too much about how the battle will go. The instruction is not to stop fighting. It is to fight without the fruits in his hands, and verse 48 names the state: evenness of mind in success and defeat. Gandhi's gloss, from his own commentary, is practically a line from the choking literature: he who is always brooding over the result often loses nerve in the performance of his duty.

The Taoists had the sharpest image. In the Zhuangzi, an archer shooting for nothing has all his skill. Shooting for a brass buckle, he is already nervous. Shooting for gold, he goes blind. "He thinks more of winning than of shooting, and the need to win drains him of power" (Zhuangzi, ch. 19, Watson translation). That is Beilock's golfers, twenty-three centuries early. Wu wei, the Taoist term usually translated as non-action, is closer to effortless action: absorbed, unselfconscious effectiveness, the cook whose knife stays sharp for nineteen years because he follows the joints instead of forcing them. Edward Slingerland, who has spent a career on the concept, points out the paradox the ancients already saw: you cannot try to be spontaneous (Slingerland, Trying Not to Try). The Confucians answered with training until it becomes second nature. The Taoists answered with letting go.

What the three agree on is the thing this essay is about. Keep the effort. Keep the interest. Drop the stake in the result. And all three are explicit that the opposite error, not acting, not caring about the work, is a failure and not a virtue. The Gita forbids inaction in the same breath. The Stoic sails. The archer still shoots.

How to borrow the stance for the things you love

You cannot decide to stop caring about your own project, and you should not want to. What you can do is manufacture the conditions that made the channels easy: distance, pre-commitment, and a judge with no stake. Here is what has worked for me, each tied to the evidence that says why.

Talk about it in the third person. Grossmann's subjects reasoned better about their own problems the moment they were asked to view them as an outsider, and a month of third-person diary writing raised the quality of their judgment. Moser's lab found the same move cut emotional reactivity with no extra mental effort (Moser et al., 2017). So when I review the game now, I write the notes as if reviewing someone else's game. "He kept the dragon art because he likes it. The data says players skip that screen." It feels silly for about a week.

Set the rules before you care. The trading engine works because the stop-loss was written down before the position existed. Every rule in it was chosen on a day when no money was on the line, and the agent's job is to refuse to change a number mid-trade. Do the same with a launch: write the success metric and the kill criterion before you build, in a sentence a stranger could check. The escalation literature says the thing most likely to keep you pouring in resources is ego threat, and a rule written in advance is the only argument ego cannot win.

Run tests you are indifferent to. The Microsoft lesson is not that most ideas fail; it is that you cannot tell which, so stop trying to. Two versions, one change, let the numbers pick. If you find yourself hoping, you have already chosen a side; test anyway and let the hope lose.

Ship the version you would ship for a client. I have never missed a client deadline to polish something the client did not ask for. I have missed my own. The question that fixes it is the same one the Gita asks: is this effort toward the work, or toward the fruit? Polish that improves what the player experiences is work. Polish that protects how I feel about it is fruit.

Give the outside view to something that does not care, and set it up to disagree. This is where the AI earns its place, with one large caveat. Left alone, a model will flatter you: in a 2026 Science study, people who described a real conflict to an affirming AI came away about 25% more convinced they were right and 10% less willing to repair (Cheng et al., 2026). So the instruction matters. Mine is standing: tell me where I am wrong before you tell me where I am right, and never confidently agree without checking first. The night it told me my own trading engine lost to an index fund was the night it became useful.

Keep the interest, drop the identity. This is the whole distinction, and the passion research draws it precisely. Harmonious passion, where you love the activity, predicts well-being and persistence. Obsessive passion, where the activity is your worth, predicts distress and nothing else. The channels

worked because my identity was never in them. The fix for the game is not to love it less. It is to stop needing it to prove anything about me.

The objections I take seriously

"You got lucky once and built a theory around it." Partly fair. The channels are one case, the platforms were having a good month for small accounts, and I cannot prove the detachment caused the reach. What I can say is that the mechanism is not mine: the self-other gap in judgment, the ego-threat driver of escalation, the founder's blind spot about his own venture and the two-thirds failure rate of expert ideas are all measured on thousands of people. The case is an illustration. The literature is the argument.

"Passion is why anyone finishes anything." Also fair, and the research agrees up to a point. Harmonious passion predicts persistence and protects against burnout. Intrinsic motivation predicts the quality of work across 183 studies, and external rewards can undermine it (Deci, Koestner and Ryan, 1999; Cerasoli, Nicklin and Ford, 2014). So the essay is not asking you to care less about the work. It is asking you to care less about the verdict. Those come apart cleanly in the data: the passion that helps is the kind you control; the passion that hurts is the kind wired to your worth.

"Not caring is how burnout looks from the inside." True, and it is the most important objection. Cynical detachment is one of the three dimensions of burnout in Maslach's model, a distancing from the work and the people in it that follows exhaustion (Maslach and Leiter, 2016). The healthy form has an older name from medicine: detached concern, enough distance to think and enough care to act. The difference is the direction. Burnout withdraws from the work to protect the self. What I am describing withdraws the self from the verdict to protect the work.

"The numbers can lie, and a person who only follows numbers will follow them off a cliff." Yes. Completion rate rewards the cheap trick as readily as the good clip, and a company that optimizes every A/B test can optimize itself into something nobody wanted. Not caring about the answer is not the same as not caring about the question. Someone with taste still has to decide what is worth testing, and that judgment is exactly what the rest of this essay says to protect from the ego, not to hand to a dashboard. The trading engine beat my record and lost to the index; a pure numbers person would have shipped it anyway.

"Detachment is a luxury of people with nothing on the line." The loss-aversion data says the bias grows with stakes, so the person with rent due is the one least able to be cool about it. I do not have a clean answer except that this is precisely why the rules have to be written before the stakes arrive, and why the hardest version of the practice belongs to the people who need it most. The Stoics were writing for slaves and emperors alike, and the Gita is addressed to a man on a battlefield.

"Emotion is information. The lesion patients who felt nothing made worse lives." Correct, and worth stating plainly. In a famous investment game, patients who could not feel out-invested healthy controls, because they did not flinch after losses (Shiv et al., 2005). The same lab's other finding is that those patients do badly at life. Emotion is a useful signal and a bad decision rule. The practice here is to read the signal and then decide by the rule you wrote when you were calm.

What to do on Monday

1. Pick the project you care about most. Write one paragraph about it in the third person, as if describing a stranger's project to a friend. Read it back. The sentence that makes you wince is the decision you have been avoiding.
2. Before the next launch, feature or trade, write down two things on a day when nothing is at stake: what success looks like in a number, and what makes you stop. Put the date on it. Do not edit it afterward.
3. Take the thing you have been defending and run it against one alternative you do not like. Let the numbers pick. If the alternative wins, ship it the same day, before the hope grows back.
4. Find your uncaring judge. A person who does not love you, a rule written in advance, or an AI with a standing instruction to tell you where you are wrong first. Ask it before every decision that feels important, especially the ones that feel obvious.
5. Practice indifference on something you hate first. If there is a channel, a chore or a format you have no feelings about, run it purely by the data for a month. Watch how good your judgment is when nothing is yours. Then carry that version of yourself back to the work you love.

The point of all of this is not to care less. It is to stop asking the work to tell you who you are, so that you can see it clearly enough to make it good.

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 I say so; where it rests on a meta-analysis I give the number of studies. A few things that are commonly said were left out on purpose because the evidence does not carry them: the Yerkes-Dodson inverted U (twelve mice, one experiment in three, failed replications), the precise TikTok hook statistics that circulate on marketing blogs (no traceable source), the DALBAR investor behavior gap (its method confuses cash-flow timing with behavior), and the claim that a financial advisor adds three points a year (a vendor estimate, not a measurement). My own channel numbers are a case of one and are presented as such.

The research brief behind this essay, with a strength-of-evidence rating on every item and the list of claims cut, runs to about 6,000 words and is available on request. The first essay in this series is The Ego Death of Society.

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, the instruction was the point.