

Fragility and the Risk Paradox

The Question

What is the opposite of fragile?

I first came across this seemingly simple question while reading a book by Nassim Taleb, and it changed the way I look at people, organisations, and the systems that support them.

While the question may seem simple, the implications of the answer (which probably isn't what you think it is) are quite complex. A discussion about fragility is also not for everyone. In my experience it's quite normal for people to initially resist the concepts, and to need time to process them. So while my intention is not to annoy you to the point where you stop reading, I do want to challenge some of the traditions that are no longer serving our industry.

The Answer

Did you say that the opposite of fragile is strong, or resilient, or robust? I know that's where I went when I first read the question, and that's where were most people go when I discuss this topic in open forums. The problem however is that the opposite of fragile is not robust, and let's now discuss why.

Fragility and Robustness

First, if we stress something that is fragile it breaks (or is harmed), and usually in a pretty permanent way. Glass is an obvious example in that when it breaks there is no coming back, and the breaking (stress) point is relatively low. Second, if we stress something that is robust then nothing happens (up to a point of course). Concrete, steel, rubber and bricks are all examples of things that are robust because we what we rely on is their ability to be stressed and remain essentially the same.

But, here is the next key question. Are there things that gain or improve from stress?

It could look like this:

Stress	Stress	Stress
∨	∨	∨
Fragile	Robust	?????????
∨	∨	∨
Harm / Damage	Nothing / No Change	Gain / Improvement

The Opposite of Fragile

As soon as we slow down the discussion and point it out it seems obvious right? Of course there are things that can gain from stress. One of the most obvious is the human body. We know that to strengthen it we need to put it under some stress, and we know that removing all stress is unhealthy. In fact most organic systems fit this description, and depending on your personal beliefs, it is essentially describing the process of evolution.

However there is no word for it (in English anyway), and so Taleb coined the term antifragility. Antifragile things gain from disorder. They become stronger following stress. For things that are antifragile, stressors are cues or information, not problems.

The Next Big Question - So What?

So far all we have done is come up with a word for something that already exists. But what if I now asked- How would you like your organisation and systems to be? I'm yet to meet a leader who doesn't say, antifragile.

Interestingly, I've also met plenty that assume their organisations are already antifragile. In most cases that are not, so this is again a reason why I believe there is a need look at this a little closer because the lack of awareness of the discrepancy is important.

At this stage I also think it's important to understand that I'm not suggesting that every organisation should be antifragile. I'm merely pointing out that when asked all leaders say want it, few actually understand it, and almost all are unaware that many things we do in traditional safety and risk actually have the effect of adding fragility and risk.

The Paradox

That many of the orthodox safety and risk strategies and interventions currently employed within organisations are actually adding fragility, and therefore risk, is the primary paradox of fragility. Typically, adding fragility and risk in this way happens over time and behind the scenes. It's one of the hidden trade-offs that are always there, but rarely looked for. It's not obvious until something goes wrong, but when there are unexpected outcomes, fragile organisations are more likely to take a big. Antifragile organisations are more able to not only keep operating, but then come out the other side in a better position. Antifragile organisations learn from these events, rather than just sending out alerts and hoping they never happen again.

What Makes Things More Or Less Antifragile?

As mentioned earlier, this is a complex and broad topic. There are numerous factors that influence fragility, some of which are about mindset and culture, and others that are more structural and systemic. In providing just an introduction to this topic, my hope that it triggers an interest to investigate some of the detail further, because there is a lot of detail available.

Antifragile organisations:

- Treat **risk as non-linear**. They know that some risk can be OK, and others can be disastrous. Fragile organisations treat all risk as unwanted (think zero harm). They seek to ban all manual handling, rather than understanding that some manual effort is good for the body, but too much can be harmful. It's also how we end up banning all ladders from a site, or issuing warnings to people who don't hold the handrail on stairs, while ignoring underlying cultural issues that are contributing to adverse mental health issues. Antifragile organisations know risk is a wicked problem and consider context and circumstances, fragile organisations think that all risk is bad.
- Understand that **stressors and unexpected events are valuable sources of feedback** about their systems and processes, rather than things to be avoided at all cost. Fragile organisations are scared of unexpected events, and this is seen in how people react to incidents. Antifragile organisations are scared

when there are no unexpected events, because this means there is less information coming in, and that it is most likely being hidden.

- Place more value on **learning that focuses on developing understanding**, rather than just transferring content. Fragile organisations have a false sense of the competence of their people because everyone passes the test. Antifragile organisations understand that know how takes time to develop. That on the job/ experiential training is the most effective way for people to learn. They still have inductions, but they don't assume that that is where people learn how to be safe.
- Resist tightly coupled, centrally controlled structures. They have **flatter, more distributed control structures** that are better able to learn and are more tolerant to failures. Importantly, they understand and accept the trade-offs that come with that approach. The more distributed the system is, the less visibility and control there is at the top. Allowing parts of the business to discern and learn on their own is antifragile, but of course impacts control and command. But it allows one part to make a mistake that doesn't sink the whole business, yet can be learnt by everyone. Fragile organisations cling to control from the top, which impacts learning and trust.
- Have a policy where they **gather information first, and then decide if an intervention is required** following an unexpected event. Fragile organisations tend to have a policy of intervention. They typically are not even aware of it because it is so entrenched. Just consider the reply when you ask why incidents are investigated. Is the answer, "to find a cause to prevent recurrence"? Can you see how the decision to intervene (fix) is already decided? This is really awkward then when we investigate and genuinely don't know what happened. It makes sense then that the only think left to fix is the person. In fragile organisations there is very little space to say "we just don't know what happened". Antifragile organisations investigate first, and then decide if further action is needed.
- Take a **longer-term view**, which leads to more of a "wait and see" approach, rather than reacting to every event. Fragile organisations tend to react to every unexpected outcome. Consider how some organisations react to changes in LTIFR. If it dips down it's celebrated as an indicator that things are safer, and of course if it rises up it's time for a change in direction and strategy. This reaction to short-term issues adds fragility through uncertainty and destabilisation, and reinforces incorrect messages, such as suggesting that LTIFR actually means something.
- **Expect that people deviate** from established processes as a normal thing, rather than expecting 100% compliance to these processes. Just think about driving and whether we are 100% compliant the whole time, and what the road system would look like if it was designed that way. So Antifragile organisations expect variation. It gets designed into systems and processes, and therefore builds natural tolerance. Small errors are less likely to lead to negative outcomes because they can be detected and corrected. In a system that relies on compliance, small errors are accentuated into sever outcomes more often.

Quick Recap

OK, so where are we now? The opposite of fragile is not robust, it's antifragile. Antifragile things gain from disorder, and nearly everyone wants their systems and organisations to be antifragile. However it turns out that some of the things we do to reduce risk might actually add fragility and risk.

Fragility is Fine, As Long As You Know

Once again, I'm not suggesting that every organisation needs to be antifragile. I would suggest however that it's important to understand the implications of some of the things we do in safety and risk. For example, implementing a new process to encourage safety conversations within your organisation is a great idea. However nothing happens in isolation, and so setting a KPI, and then requiring a form to be filled out every time, will typically add fragility.

The competing demands (have conversations, meet KPIs, and complete the form) create trade-offs over time, and in all but the most mature culture the quality will lose out to quantity at critical times. It applies a centralized control structure to an organic process (people meeting). There is also fragility introduced because of the meaning attributed by the organisation to the numbers. "We are safer because we are hitting our KPIs for conversations". In this statement there is no discussion of meaning or quality because the structure of measurement wins out of the process of meeting. Without the KPIs in place there is of course no objective way to know if they are being done (no visibility), except to go and engage with people and check. This is more meaningful, but takes time and effort and is less measurable. It's complex (wicked), and there is no easy answer, but surely it's still worth the conversation.

Closer to Antifragile

At a broad level then, there are some things that can be done to move closer to antifragile, if you want to.

Antifragile Organisations:

- Ensure decisions are being made or influenced by people that will be impacted by them (**skin in the game**). In short, whether you are going to be impacted by a decision, or not, influences the decision making, particularly over time. Changing a shift roster without involving anyone working those shifts is more likely to introduce fragility. Taleb says it introduces tail risk, or uneven distribution of risk, where people are making decisions, but are not directly exposed to the outcomes of those decisions.
- **Think about the ethics** of an intervention (why they are doing something and how it will be used) before they implement. For example, was LTIFR originally meant to be used to do things like award bonuses or assess the suitability of contractors? In reality, LTIFR is used this way, and it introduces fragility because people fudge the numbers to meet the expectation. It means the original purpose (feedback on safety) is no longer in place. Antifragile organisations work out why they are going to do something before they do it to try and prevent misuse later. They might say that they are going to implement "golden rules", but only for education purposes, and that they are never intended to be used for discipline. This approach would be more likely to encourage reporting of near misses, but drive them underground.
- **Separate investigation and intervention.** As mentioned earlier, they have a protocol to determine intervention, not a policy of intervention. They investigate, gather the information, and then determine if they need to do something (not what to do as that still implies something needs to be done). It's entirely possible that some investigations could have no recommendations. It could look like this, "following our investigation we believe that the organisation was taking all reasonably practicable measure to prevent the incident and therefore that are no correctly actions required".

- **Have smaller, flatter, more loosely coupled** systems and organisational structures, and are OK with the trade-offs that come with that. They consider that engagement and connection is more important than data and forms. They think trust is important, and also understand that asking for proof damages trust.
- Try to **maintain optionality**, rather than always seeking to limit choice. Fewer options equals more fragility. Again, it isn't about giving free choice to all, it's about understanding the implications. So antifragile organisations might focus on an outcome (safe operations of a vehicle), but not try and tell people the only way they are allowed to achieve it, especially in circumstances where there is the potential for natural variability to be introduced. The question "what is 3 plus 3?" is fragile. It's either right, or totally wrong, but in this situation that's OK because that is the nature of maths. But when the question is "what is the safe way to do this job?", the answer may not be so black and white, so the procedure shouldn't be either. Having a procedure that dictates how to do a task with lots of variation adds fragility. Antifragile organisations accept and embrace this, and focus on ensuring clarity of the outcome, not necessarily how to get there.

Final Word

Antifragility is not a new concept, all I have done (through explaining Taleb's approach) is to give it a label. It's also completely normal that people want their systems to be antifragile. The trouble starts when we discuss what makes things more or less fragile, and what we can do about it if we choose to. This is where the concepts and strategies of antifragility require a move away from tradition. Again, it's not that everything should be antifragile, it's that we are often blind to the impact of our actions on fragility and risk.

In the end, the illusion of control through vast and complex systems and processes is an illusion, and it is a dangerous one because it adds risk in the background while giving the impression that risk is being managed in the foreground. Large, irrelevant, imposed systems shift power away from people, who in the end are the source of safety and not a problem to be fixed or controlled. It is a paradox, and it's normal that any discussion of antifragility creates discomfort. But isn't the thought of just doing what we are doing now, forever, even more uncomfortable?

Dave Whitefield
Director
People and Risk
0412 782 191