INSIGHTS · GENERATING INSIGHT

Failure Analysis and Post-Mortem Reviews as Sources of Corrective Insights

How organisations can learn from failures through structured post-mortem reviews, covering why failure is hard to learn from, hindsight and outcome bias, systems thinking about error, psychological safety, the evidence on after-action reviews and a practical review process.

ModQu research article · 2,001 words · 21 references (APA 7th edition)

Failures contain information that successes often do not. When a product launch falls short, a campaign misses its target or a forecast proves badly wrong, the gap between what was expected and what happened reveals something about the assumptions behind the decision. Organisations that examine failures carefully can correct those assumptions and avoid repeating the same mistakes. Yet many organisations find it difficult to learn from failure. Reviews are skipped, blame is assigned and the lessons that are recorded are rarely acted on. This article explains why learning from failure is difficult, what research says about doing it well and how insight teams can use post-mortem reviews as a source of corrective insight.

Why failure is valuable

Sitkin (1992) argued that organisations learn more from small failures than from success. Success tends to produce complacency, reinforce existing routines and reduce the search for new information. Failure, by contrast, draws attention to problems, prompts people to search for explanations and makes it more acceptable to change course. Sitkin proposed a "strategy of small losses": deliberately encouraging small, manageable failures that provide learning without threatening the organisation.

Madsen and Desai (2010) tested this idea in a study of orbital launch attempts across the global space industry. They found that organisations learned more effectively from failures than from successes, and that knowledge gained from failure decayed more slowly. Larger failures produced more learning than small ones.

Learning only from success can also be misleading. Denrell (2003) showed that when people study only successful firms, they can draw false conclusions about which practices lead to success, because the firms that followed the same practices and failed are no longer visible. Risky strategies, for example, may appear more successful than they are because the many firms that tried them and failed have disappeared from view. Studying failures helps correct this bias.

Why failure is hard to learn from

Cannon and Edmondson (2005) identified several barriers to learning from failure in organisations. Technical barriers include the difficulty of detecting failures, especially small ones, and of analysing their causes. Social barriers include the reluctance of people to report failures that might reflect badly on them and the tendency of managers to punish the bearers of bad news. Many organisations, they found, focused on large, visible failures and missed the small ones that could have provided early warnings.

Tucker and Edmondson (2003) studied how hospital nurses responded to operational failures, such as missing supplies or incorrect information. In most cases, nurses worked around the problem and continued with their tasks, but did not report it or try to change the system that caused it. This "first-order problem solving" kept work going but meant that the same problems recurred. Learning required "second-order problem solving": reporting the problem and addressing its root causes.

Near misses and the normalisation of deviance

Some of the most useful failures to study are the ones that almost happened. Dillon and Tinsley (2008) found that people who experienced a near miss, an event in which a failure was avoided partly by chance, tended to judge the underlying risk as lower and to make riskier decisions afterwards. The near miss was interpreted as evidence that the system worked rather than as a warning that it had nearly failed.

Vaughan (1996) studied the decision to launch the space shuttle Challenger in 1986. Vaughan found that engineers and managers had repeatedly observed damage to the seals that later failed, and that each time the damage was within limits they came to see it as acceptable. Over time, a pattern of small deviations from the original design expectations became the normal state of affairs. Vaughan called this the normalisation of deviance. In business, the same process occurs when slipping data quality, rising complaint rates or repeated small forecast errors are accepted as normal because nothing serious has yet gone wrong. Reviewing near misses and small deviations, not only major failures, helps insight teams detect such drift before it produces a larger failure.

Biases that distort post-mortems

Reviewing a failure after the fact is subject to predictable biases.

Hindsight bias is the tendency to see past events as having been more predictable than they were. Fischhoff (1975) found that people who were told the outcome of an event judged that outcome to have been more likely than people who were not told, and that they were largely unaware of this effect. Roese and Vohs (2012) reviewed later research and described how hindsight bias leads people to believe they "knew it all along", which makes failures seem like obvious mistakes rather than reasonable decisions that turned out badly.

Outcome bias is a related problem. Baron and Hershey (1988) found that people rated the quality of a decision more highly when it had a good outcome than when it had a bad one, even when the information available at the time of the decision was identical. In a post-mortem, this bias leads reviewers to condemn decisions that were reasonable given what was known, and to praise decisions that turned out well by chance.

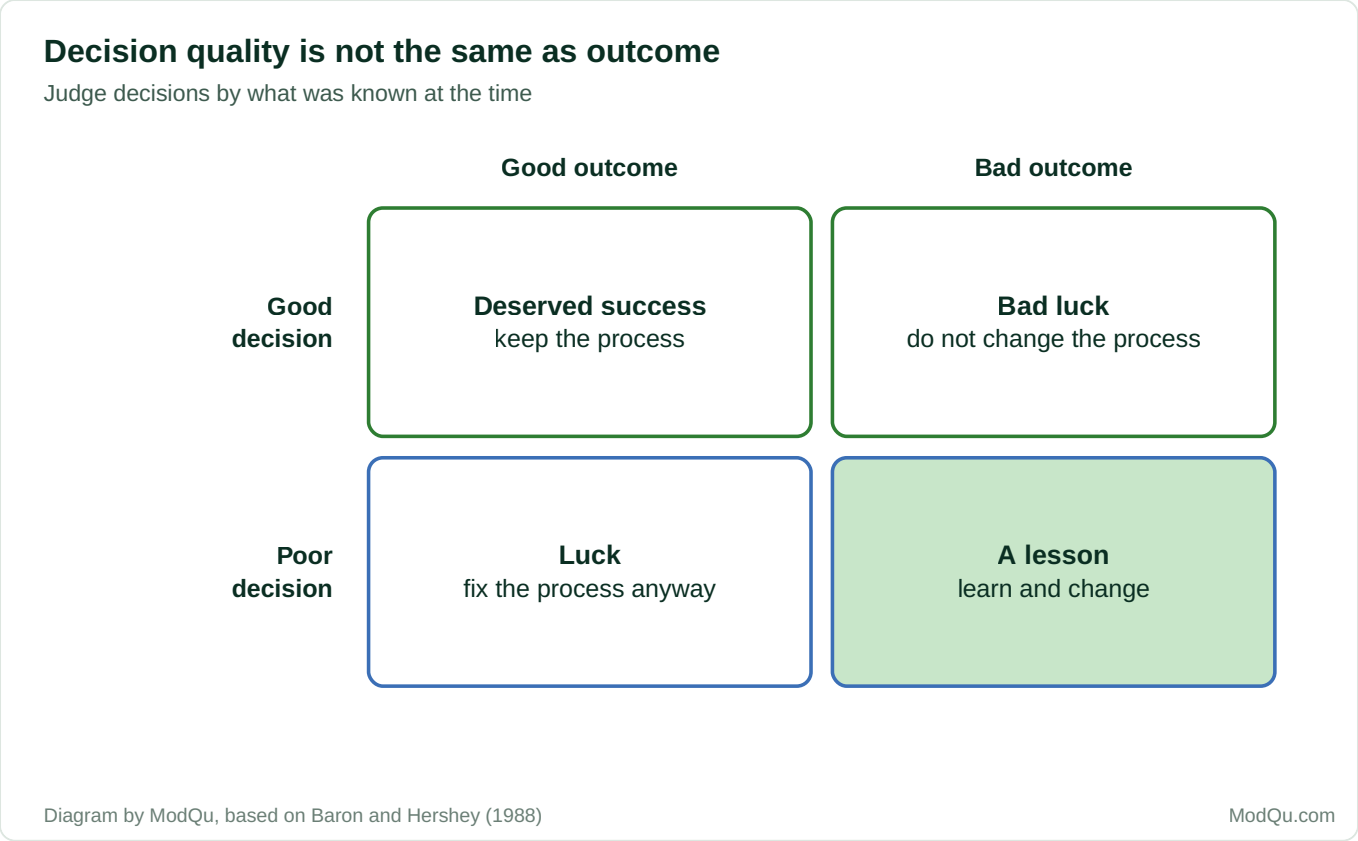


Figure 1. Separating the quality of a decision from its outcome. Diagram by ModQu, based on Baron and Hershey (1988).

The practical lesson is that a post-mortem should judge decisions by the information and options available at the time, not by what happened afterwards. One useful question is: given what we knew then, would a reasonable person have made the same decision?

From blaming people to understanding systems

Reason (1990, 2000) distinguished between two approaches to human error. The person approach focuses on the errors of individuals and attributes them to carelessness, inattention or poor motivation; it responds with blame, retraining and discipline. The system approach starts from the assumption that people will make errors and asks why the system allowed an error to cause harm. Reason described how failures in complex systems usually arise when several weaknesses in defences line up, so that a single error passes through all of them. Reason argued that the system approach is more effective, because blaming individuals leaves the weaknesses in place for the next person.

Dekker (2014) argued that what looks like human error after the fact usually made sense to the people involved at the time, given their goals, knowledge and the pressures they faced. The purpose of a review is to understand why their actions made sense, because that reveals the conditions that will lead others to make the same choices.

For insight work, the system approach changes the questions a post-mortem asks. Instead of "who made the wrong forecast?", it asks "what information, methods and incentives led to a forecast that was wrong, and would they lead someone else to the same result?"

Psychological safety

Honest reviews require people to speak openly about their mistakes. Edmondson (1999) found that teams with high psychological safety, in which members believed they would not be punished or humiliated for raising problems or admitting errors, engaged in more learning behaviour. Edmondson (2011) argued that organisations should distinguish between different kinds of failure. Preventable failures in routine work usually reflect deviations from known procedures. Complex failures arise from novel combinations of factors in complicated systems. Intelligent failures occur in experiments and new ventures at the frontier of knowledge, where failure is a necessary part of discovery. Treating all failures as blameworthy discourages the reporting and experimentation that organisations need.

What the evidence says about reviews

Structured reviews of past events, known as after-action reviews, after-event reviews or debriefs, have been studied in military, medical and business settings. Ellis and Davidi (2005) studied soldiers undergoing navigation training and found that those who conducted after-event reviews of both successful and failed exercises improved more than those who did not. Reviewing successes as well as failures was valuable, because it helped people understand why things went well and not only why they went badly.

Tannenbaum and Cerasoli (2013) conducted a meta-analysis of debriefs and found that teams and individuals who conducted well-structured debriefs improved their performance by about 25% on average compared with control groups. Debriefs were most effective when they were facilitated, structured and aligned with specific objectives.

Keith and Frese (2008) found in a meta-analysis that training which encouraged people to make errors and learn from them, known as error management training, led to better performance on new tasks than training that tried to prevent errors. Errors provided feedback that helped people build a deeper understanding of the task.

A process for post-mortem reviews

A useful post-mortem review has five stages.

The first is to establish what happened, in sequence, using records rather than memory where possible. Reconstructing the timeline before discussing causes reduces the influence of hindsight.

The second is to compare what happened with what was expected. What assumptions did the original plan rest on, and which of them proved wrong?

The third is to identify contributing causes, looking at systems, information and incentives as well as individual actions. Tools such as cause-and-effect diagrams (Ishikawa, 1976) help a team consider several categories of cause, such as people, methods, data and environment, rather than settling on the first explanation.

The fourth is to draw lessons and decide on specific changes, each with an owner and a deadline.

The fifth is to follow up, checking whether the changes were made and whether they worked.

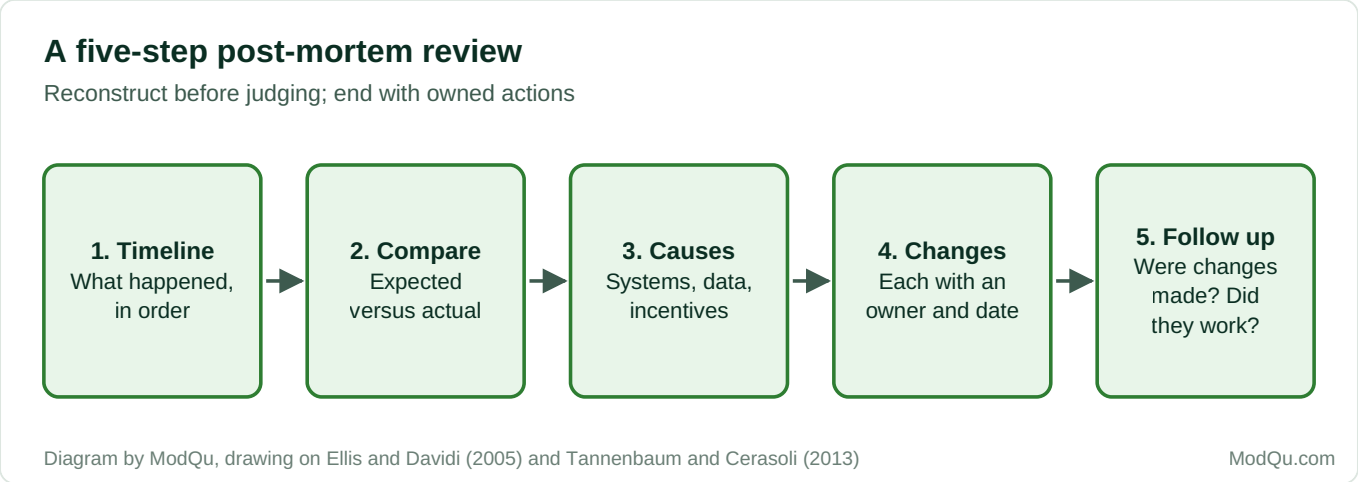


Figure 2. A five-step post-mortem review. Diagram by ModQu, drawing on Ellis and Davidi (2005) and Tannenbaum and Cerasoli (2013).

Looking forward: the pre-mortem

Some of the benefits of a post-mortem can be obtained before a project starts. Klein (2007) proposed the pre-mortem: before a decision is finalised, the team imagines that the project has failed and writes down the reasons why. Mitchell et al. (1989) found that imagining that an event had already occurred increased people's ability to identify reasons for its outcome. A pre-mortem uses this effect to bring hidden risks to the surface while there is still time to address them.

An example

A consumer goods company launched a new drink that sold less than half its forecast volume in the first six months. An initial discussion blamed the sales team for poor execution. The insights team proposed a structured review.

The review reconstructed the timeline from the original concept tests to the launch. It found that the concept tests had used a description that emphasised health benefits, while the final packaging emphasised taste. The forecast had been built on concept test scores without adjusting for the typical gap between stated intentions and actual purchases. The launch had also been delayed by two months, so that it reached shops after the summer peak. None of these causes was the fault of the sales team.

The review produced three changes: concept tests would use the final packaging and positioning; forecasts would be adjusted using the company's own record of past launches; and launch dates would be reviewed against seasonal demand. The insights team followed up a year later and found that forecast errors for subsequent launches had fallen.

Common mistakes

Common mistakes in post-mortems include holding them only for large failures, which misses the smaller failures that provide early warnings; focusing on who was at fault; reviewing too long after the event, when memories have faded and records are incomplete; producing long lists of lessons without specific actions; and never checking whether recommended changes were made. Another mistake is reviewing only failures, which misses the opportunity to learn why successful projects succeeded and whether their success depended on luck.

Conclusion

Failures are among the richest sources of corrective insight available to an organisation, but only if they are examined carefully. Research shows that organisations can learn more from failure than from success (Madsen & Desai, 2010), that structured reviews improve performance (Tannenbaum & Cerasoli, 2013) and that hindsight and outcome biases can distort what is learned (Baron & Hershey, 1988; Fischhoff, 1975). Effective reviews reconstruct events before judging them, focus on systems rather than individuals, take place in a climate of psychological safety and lead to specific changes that are followed through.

References

Baron, J., & Hershey, J. C. (1988). Outcome bias in decision evaluation. *Journal of Personality and Social Psychology, 54*(4), 569–579. https://doi.org/10.1037/0022-3514.54.4.569

Cannon, M. D., & Edmondson, A. C. (2005). Failing to learn and learning to fail (intelligently): How great organizations put failure to work to innovate and improve. *Long Range Planning, 38*(3), 299–319. https://doi.org/10.1016/j.lrp.2005.04.005

Dekker, S. (2014). *The field guide to understanding 'human error'* (3rd ed.). Ashgate.

Denrell, J. (2003). Vicarious learning, undersampling of failure, and the myths of management. *Organization Science, 14*(3), 227–243. https://doi.org/10.1287/orsc.14.3.227.15164

Dillon, R. L., & Tinsley, C. H. (2008). How near-misses influence decision making under risk: A missed opportunity for learning. *Management Science, 54*(8), 1425–1440. https://doi.org/10.1287/mnsc.1080.0869

Edmondson, A. (1999). Psychological safety and learning behavior in work teams. *Administrative Science Quarterly, 44*(2), 350–383. https://doi.org/10.2307/2666999

Edmondson, A. C. (2011). Strategies for learning from failure. *Harvard Business Review, 89*(4), 48–55.

Ellis, S., & Davidi, I. (2005). After-event reviews: Drawing lessons from successful and failed experience. *Journal of Applied Psychology, 90*(5), 857–871. https://doi.org/10.1037/0021-9010.90.5.857

Fischhoff, B. (1975). Hindsight is not equal to foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance*, *1*(3), 288–299. https://doi.org/10.1037/0096-1523.1.3.288

Ishikawa, K. (1976). *Guide to quality control*. Asian Productivity Organization.

Keith, N., & Frese, M. (2008). Effectiveness of error management training: A meta-analysis. *Journal of Applied Psychology*, *93*(1), 59–69. https://doi.org/10.1037/0021-9010.93.1.59

Klein, G. (2007). Performing a project premortem. *Harvard Business Review*, *85*(9), 18–19.

Madsen, P. M., & Desai, V. (2010). Failing to learn? The effects of failure and success on organizational learning in the global orbital launch vehicle industry. *Academy of Management Journal*, *53*(3), 451–476. https://doi.org/10.5465/amj.2010.51467631

Mitchell, D. J., Russo, J. E., & Pennington, N. (1989). Back to the future: Temporal perspective in the explanation of events. *Journal of Behavioral Decision Making*, *2*(1), 25–38. https://doi.org/10.1002/bdm.3960020103

Reason, J. (1990). *Human error*. Cambridge University Press.

Reason, J. (2000). Human error: Models and management. *BMJ*, *320*(7237), 768–770. https://doi.org/10.1136/bmj.320.7237.768

Roese, N. J., & Vohs, K. D. (2012). Hindsight bias. *Perspectives on Psychological Science*, *7*(5), 411–426. https://doi.org/10.1177/1745691612454303

Sitkin, S. B. (1992). Learning through failure: The strategy of small losses. *Research in Organizational Behavior*, *14*, 231–266.

Tannenbaum, S. I., & Cerasoli, C. P. (2013). Do team and individual debriefs enhance performance? A meta-analysis. *Human Factors*, *55*(1), 231–245. https://doi.org/10.1177/0018720812448394

Tucker, A. L., & Edmondson, A. C. (2003). Why hospitals don't learn from failures: Organizational and psychological dynamics that inhibit system change. *California Management Review*, *45*(2), 55–72. https://doi.org/10.2307/41166165

Vaughan, D. (1996). *The Challenger launch decision: Risky technology, culture, and deviance at NASA*. University of Chicago Press.