CONDITIONS FOR LEARNING

Learning is a critical to resilient performance, but we tend to focus most safety learning on things that have gone wrong. In this Op Ed, **Erik Hollnagel** argues that everyday work provides the best basis for learning. To achieve this, three conditions for learning are proposed.

It is practically an axiom of safety management that learning must be based on things that have gone wrong – from accidents and incidents (Kletz, 2001, originally 1988). There is an intuitive logic to this. Things that go wrong are unwanted and costly both in the short and the long run and may also disrupt planned operations. It therefore stands to reason that failures should be prevented as far as possible and it is assumed that this is best achieved by analysing them to find their underlying causes. Once these have been found then steps can be taken to ensure that they will not happen again, at least according to the current safety dogmas.

The purpose of learning is generally accepted to be the acquisition of knowledge or skills through study, experience, or being taught. The ultimate purpose of learning is, of course, to change behaviour – either so that something can be done better, so that something different or new can be done, or so that certain specific behaviours can be avoided. Psychology, in fact, refers to a special form of learning called avoidance learning, defined as the process by which an individual learns a behaviour or response to avoid a stressful or unpleasant situation. Learning that does not lead to a change in behaviour is of very limited interest and also very hard to verify.

If we look beyond safety management and consider learning in general, for instance as one of the four potentials that are proposed as the basis for resilient performance (Hollnagel, 2018), then it can be argued that three conditions are necessary for learning to take place.

Condition #1: The Opportunity to Learn

In order to learn anything, there must clearly be an opportunity to learn. This condition is so obvious that it usually is just taken for granted. If nothing happens then nothing can be analysed and nothing can be learned. But this condition actually creates a small paradox for learning to improve safety. If learning is reserved for situations where something has gone wrong – for accidents and incidents – then there will actually be few opportunities to learn. The (relative) absence of accidents is, of course, desirable for an organisation or a business, but it is not a good basis for learning. Efforts to improve safety therefore have the unintended and undesirable side-effect that they reduce the opportunity to learn, the extreme case being the zero accident principle (Zwetsloot et al., 2013). This evokes the so-called 'fundamental regulator paradox'.

"The task of a regulator is to eliminate variation, but this variation is the ultimate source of information about the quality of its work. Therefore, the better the job a regulator does the less information it gets about how to improve." (Weinberg and Weinberg, 1979, p. 250).

In relation to safety, this means that if something rarely or never happens, then it is impossible to know how well it works – and probably also difficult to justify investments to further improve how the system works.

For the sake of learning it would make sense to increase the opportunity to learn, which means increasing the number of conditions where something can be learned. In relation to safety, this appears to constitute a paradox, for who would seriously propose that we should increase the number of accidents? But the paradox only exists as long as safety is defined as the freedom from accidents and incidents. If instead of trying to learn what not to do we tried to learn what to do, then the focus would change from situations where things went wrong to situations where things went well. The purpose of learning would likewise change from trying to avoid something to trying to approach something. This corresponds to a definition of safety as a condition where as much as possible goes well, also referred to as 'Safety-II' (Hollnagel, 2014).

In the daily life of an organisation nearly everything goes well, not in the sense that there is a perfect agreement between performance and rules, procedures, and regulations, but in the sense that the outcomes are acceptable to the organisation itself as well as to its customers. (This, by the way, applies not only to safety but also to quality, productivity, etc.) Indeed, the less spectacular the outcomes are, the more opportunities for learning there will be, both in the sense that there will be more situations to learn from and in the sense that the cost of making the necessary changes will be significantly lower than for classical accident prevention.

Condition #2: The Similarity Between Situations

Learning cannot take place unless there is enough similarity between situations to make generalisation possible. This allows people and organisations to recognise situations based on their experience, and therefore to respond more efficiently. Without having some

patterns or traits to look for, each situation would have to be analysed anew, which is neither reasonable nor practicable.

If there are few or no similarities between situations then the nature of learning changes from generalising across situations to become the ability to remember an ever-growing set of individual situations. This corresponds to a kind of signature-based detection that was the initial approach used by anti-virus software. However, as the number of signatures grow, so does the time it takes to make a thorough search of them – and not being thorough would defeat the purpose.

In relation to safety, accidents are usually rare events – just as they should be. This is not the best condition for learning, but to make matters worse, accidents tend to be less similar as they become more serious. This is unfortunate because of the dogma that it is more important to learn from severe accidents than from minor ones. The psychological reasons for this are obvious: the more serious and disruptive the outcomes of an event are, the higher the motivation is to ensure that it does not happen again. But it also makes general learning more difficult and leaves signature-based learning as the only option.

The situation is completely different if learning adopts a Safety-II perspective and includes things that go well. These not only happen far more frequently, but also occur as clusters of similar activities. It is therefore much easier to learn from them, to understand their characteristics and to find ways to improve or facilitate them. The clusters or patterns are furthermore not only the foundation for understanding how things go right, but also for understanding how they occasionally go wrong.

Condition #3: The Opportunity to Verify that Learning has Taken Place

The proof of the pudding is in the eating, as the proverb says. In relation to learning this means that there must be some recognisable difference in behaviour before and after learning. The

"Everyday work provides the best basis for learning while accidents provide the worst."

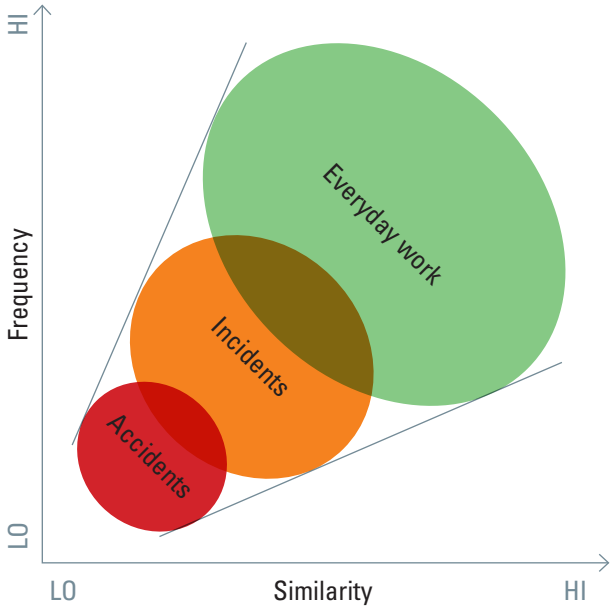

"The purpose of learning is after all to change behaviour so that certain outcomes become more likely and others less likely."

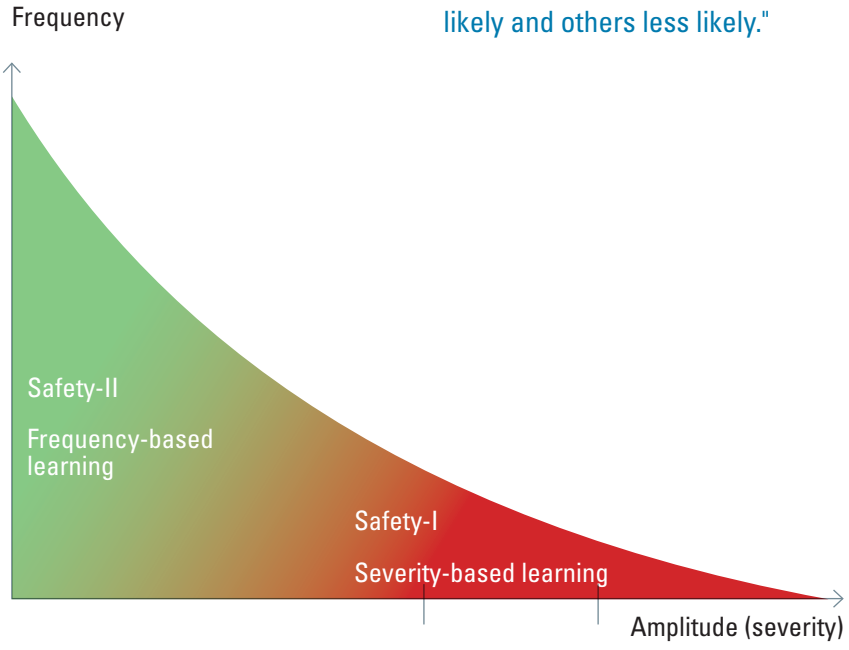

purpose of learning is after all to change behaviour so that certain outcomes become more likely and others less likely. This is notably difficult to establish for learning from accidents, since it requires that the same or a similar condition occurs again. While incidents and minor accidents may occur from time to time – but hopefully not too often – serious or severe accidents neither happen frequently nor are they very similar. The lack of feedback means that there are few opportunities to verify that the lessons that have been learned – whatever they may be – actually work as intended. This is unfortunate not least because the cost of learning from such cases can be quite high.

In contrast to that, learning from work that goes well can easily be verified. There is no need to wait for another accident to happen or even for the same accident to be repeated. Things that go well happen all the time which makes it easy to verify the effects of learning. It is, of course, not possible to look at everything that happens, but neither is it usually possible to learn from everything that goes wrong. For practical reasons a selection must be made. In the case of things that go wrong, the traditional criterion is severity of outcomes. In the case of things that go well, a good criterion would be to look at what happens most frequently. It is also easy and affordable to make adjustments and improvements, and therefore to learn incrementally and continually rather than by large jumps or steps.

Conclusions

The essence of the three conditions described above is captured in the figure below. Here, three different categories of events are shown relative to each other in terms of frequency of occurrence and severity of outcome. Accidents are low in both frequency and similarity, everyday work happens frequently – in fact it happens all the time – and has high similarity, while incidents are in between. Relative to the three conditions described above, everyday work provides the best basis for learning while accidents provide the worst. This, of course, assumes that everything happens in basically the same way and that there are no special causes of failures that appear Deux Ex Machina when something goes wrong but otherwise just lie in wait.

Accidents attract attention because they are unexpected while everyday work – things that go well – is more or less invisible. "Reliable outcomes are constant, which means there is nothing to pay attention to" (Weick, 1987). So instead of limiting learning to what we do not want to happen to find out what we should not do, it might be worthwhile also to learn from what we want to happen in order to find out what we could do.

"Learning from work that goes well can easily be verified. There is no need to wait for another accident to happen or even for the same accident to be repeated."

Erik Hollnagel is a senior professor of Patient Safety at Jönköping University (Sweden), Visiting Professorial Fellow at Macquarie University (Australia), and Visiting Fellow at the Institute for Advanced Study, Technische Universität München (Germany). Erik's professional interests include industrial safety, resilience engineering, patient safety, accident investigation, and modeling large-scale socio-technical systems. He is the author/editor of 28 books, as well as a large number of papers and book chapters.

hollnagel.erik@gmail.com

References

Hollnagel, E. (2014). *Safety-I and Safety-II: The past and future of safety management*. Farnham, UK: Ashgate.

Hollnagel, E. (2018). *Safety-II in practice*. Abingdon, UK: Routledge.

Kletz, T. (2001). *Learning from accidents* (3rd Edition). Butterworth-Heinemann Ltd.

Weick, K. E. (1987). Organizational culture as a source of high reliability. *California Management Review*, *29*(2), 112-127.

Weinberg, G. M. & Weinberg, D. (1979). *On the design of stable systems*. New York: Wiley.

Zwetsloot, G. I. J. M., Aaltonen, M., Wybo, J.-L., Saari, J., Kines, P., & Op De Beeck, R. (2013). The case for research into the zero accident vision. *Safety Science*, *58*, 41–48.