Chapter 14
High-consequence Risk with Low Likelihood

High-consequence risk with low likelihood is a challenge within risk management in general. And even more so within the domains of risk management, like cyber-security, where human intentions and behavior are important. This challenge is, however, not just one challenge, but rather a family of related challenges that should be treated separately. We distinguish between:

- Incidents that occur as complete surprises without ever having been considered; for example, something almost unthinkable that has never happened before, like al Qaeda's attack in the USA on September 11, 2001.
- Incidents whose unlikely occurrences are just likely enough to allow them to be anticipated; for example, the awakening of a volcano whose last eruption occurred 500 years ago.

In the literature, the former are commonly referred to as "black swans," while "gray swans" is used to denote the latter.

In our opinion risk assessment is of little help in identifying black swans. Risk assessment is basically a tool for obtaining a consistent picture of risk knowledge already available implicitly or explicitly within the context of the target of assessment. This knowledge is in most cases insufficient to identify black swans, but sufficient in the case of gray swans. In the following we first address black swans and then gray swans.

14.1 Dealing with Black Swans

A *black swan* is an incident that is extremely rare and unexpected, but has very significant consequences [49]. The "black swan" metaphor has historical origins. In sixteenth century Europe it was generally accepted that all swans were white. In fact, "black swan" was used in every day speech as a metaphor for something unthinkable. When sailing up a river in Western Australia January 10 1697, the members of Willem de Vlamingh's expedition were therefore highly surprised when

A. Refsdal et al., *Cyber-Risk Management*, SpringerBriefs
in Computer Science, DOI 10.1007/978-3-319-23570-7_14

they observed swans that were black [84], probably the first Europeans to do so. Tellingly they named the river “Swan River.” Nevertheless, the black swan metaphor has survived over the centuries, and recently its popularity has increased due to the writings of Nassim Nicholas Taleb [77]. Black swans are also often referred to as “unknown unknowns.”

Risk assessors are regularly confronted with questions related to the coverage of their methods, and in particular their ability to discover black swans. In our view, black swans are not likely to be discovered by risk assessment. Risk assessment is basically good for putting together and structuring a consistent picture of explicit and implicit knowledge already residing within the context of a target, not for picking up the unexpected of which there is no knowledge.

This does not mean that there is nothing we can do to prepare for black swans, only that risk assessment is not the right tool for doing it. Black swans will occur and may harm even the most risk averse organization. Hence, independent of how carefully we as risk assessors conduct our risk assessments, we must be humble and communicate to our customer (the party on whose behalf we are assessing) that although the risk picture we deliver is as good as we can make it given the available resources and information, it may be incomplete and they still need to prepare for the unexpected and plan for the unknown. Hence, risk assessment does not make contingency planning, the act of preparing and planning for major incidents and disasters, obsolete. In fact, as we see it, developing good contingency plans is the best approach to cope with black swans.

14.2 Identifying Gray Swans

A *gray swan* is an incident which has far-reaching consequences, but, unlike a black swan, can be anticipated to a certain degree [50]. If we are not careful, gray swans may also easily be overlooked in a risk assessment situation because they are not present in the documentation that we as risk assessors have gained access to. They may also be overlooked because we are not interacting with the right group of stakeholders or because we are not able to extract the required information. Although gray swans may be very unlikely to happen in the short term, they may in principle occur within the next hour, and with grave consequences. Hence, their identification is essential.

By definition, there is knowledge of all relevant gray swans within the context of the target of assessment, if not explicitly written down in some document then at least implicitly within the mind of a stakeholder or deducible from the available data. Our task as risk assessors is to extract this knowledge so that it becomes a part of our risk assessment. Whether we succeed or not depends on our approach to risk assessment and the resources available to us.

In order to uncover the true nature of some phenomenon it is often fruitful to try to observe this phenomenon from different viewpoints. This has several implications for how we conduct risk identification. One implication if we conduct the risk identification based on interviews or workshops, is that we should try to make sure that the group of people we are interacting with contains representatives for each relevant stakeholder role. After all, the perspective of a decision maker, for example, is quite different from that of an internal software developer, which again is different from that of an external consultant hired in for a period of say six months.

Another implication is that given the availability of the necessary resources and budget, we may split the risk identification into a set of independent processes each taking a slightly different approach. In one process we might, for example, start from the assets and try to identify how they may be harmed; in another process the starting point could be known threats and their potential for attacking the target; a third alternative would be to start from known vulnerabilities; and so on. A third implication is that it might be a good idea to embed the use of different kinds of tools in the risk identification process, each providing a new perspective on the target. We may, for example, use a combination of automatic vulnerability scanners, penetration tests, and monitoring tools to provide input to the risk identification.

14.3 Communicating Gray Swans

Consider a gray swan for which we assume there is a very small and exact likelihood, meaning close to no uncertainty, such as the likelihood of an attacker guessing an eight character password by sheer luck. In this case, our challenge as risk assessors boils down to the problem of communicating a very small number in such a way that its size is fully comprehended by those required to act upon it, namely decision makers.

How successful we are at this task may have great impact on the success of the decision finally made. As already explained in Sect. 12.3, we should avoid using probabilities; natural frequencies are more likely to be understood. How we present a frequency also matters a lot. For example, the frequency 0.000005 per year corresponds to the frequency 1 in 200,000 years which is the same as the frequency of once since the beginning of the human race. In general, it makes sense to present very small (and very large) numbers by relating them to entities providing a suitable perspective. On the other hand, if there is considerable uncertainty as to how small the likelihood of the gray swan is, then also this uncertainty must be communicated to relevant decision makers; for example, in terms of an upper and lower bound. In this case, there will be two small numbers to communicate. A good strategy is again to relate the numbers to some entity of quantity providing intuition, and in such a way that also the size of the uncertainty interval is fully comprehended. For example, zero to five times since the beginning of the human race.

14.4 Dealing with Gray Swans

Assume we are in the treatment phase, trying to aid the decision makers in making the right decision regarding a gray swan. If the gray swan is treatable at low cost, then it should clearly be treated. Unfortunately, this is normally not the case. For some gray swans the consequences are so grave that reducing the likelihood is the only viable alternative, even at very high cost. There are however also gray swans for which this is financially unfeasible and then the option left is to try to reduce the consequence. One obvious strategy to avoid cyber-attacks on critical infrastructure is to disconnect the critical infrastructure from cyberspace. However, in most cases this is not financially feasible. A specialized contingency plan may be a good option. A gray swan, in contrast to a black swan, is something we have knowledge or experience about in one form or another. The contingency plan for a gray swan may therefore be much more specialized and also much more effective than a contingency plan for black swans.

14.5 Recognizing Gray Swans in Cyberspace

The challenge of estimating high-consequence risks with very low likelihood is a family of challenges. What is said above regarding black and gray swans is also valid for cyber-risk. Risk assessment is, as in the general case, mainly suited to capturing gray swans. So what is a gray swan in cyberspace? As we see it, many zero-day vulnerabilities are gray swans, for example. In the same way as economists do risk assessments taking a stock exchange crash into consideration, computer scientists do risk assessments addressing zero-day vulnerabilities. We know from experience that security relevant bugs may pop up even in mature software that has been carefully tested and has been in use for many years. Since cyber-systems are computerized to a large degree we have possibilities for testing, surveillance, and monitoring that may not easily be implemented in the general case, and that may make the detection of gray swans easier. On the other hand, the existence and influence of cyberspace has a complicating effect. In particular, cyberspace is highly dynamic. Hence, a valid threat picture today may not be valid tomorrow. To cope with the dynamics of cyberspace and therefore also of cyber-systems we may aim for more dynamic risk models whose risks, vulnerabilities, and threats are defined, measured, and as much as possible updated automatically in real time as functions of low-level indicators.

14.6 Further Reading

We have already recommended the influential work [77] of Nassim Nicholas Taleb on black swans as well as Gerd Gigerenzer's book [20] on natural frequencies. Regarding the comprehension of numbers there is also specialized literature available [65]. Contingency planning is a large subject on which many have published; see for example ISO 22301 on societal security [27].

For a classification of evolution in the context of risk assessment and a corresponding case study, see [46, 48]. Initial ideas for how to automatize or semi-automatize risk assessment to keep up with scaling and system evolution have been proposed by several authors [68, 42, 69]. The field is however immature.