Risk Management Series - Part 6: Estimating Probability of Occurrence

Foreword

MEDIcept presents this ongoing series of articles focused on the implementation and practical conduct of risk management in the medical device industry to provide practitioners with insight into how to apply risk management principles and tools to improve the performance and safety of their devices; and, as an added benefit, to maintain compliance with risk management standards.

Our team at MEDIcept publishes these articles to capture best practices, to explore the more challenging aspects of maintaining risk management systems over the long term, and to elicit discussions among practitioners.

To this last point, if you have questions or comments on the issues discussed, or if you have recommendations for topics to consider in the future, please let us know: 508-231-8842.

In our last article, *Part 5: Assessing Severity,* we addressed the importance of doing your "homework" before attempting to assess the severity of identified hazards/harms, clearly defining the use environment and user profile being considered, and distinguishing between harms that occur as a direct result of the hazard/hazardous situation and harms that will only occur as the result of a series of subsequent events. This article addresses the other major element of risk, the Probability of Occurrence of Harm.

If you've missed any of the previous articles look them up on www.medicept.com/blog/

The Elements of the Probability of Occurrence of Harm

Risk is composed of two elements: Severity and the Probability of Occurrence of Harm. Unlike the assessment of Severity, which is typically based on *qualitative* criteria used to describe the harm caused to people, property or the environment; the estimation of the Probability of Occurrence of Harm is typically based on *quantitative* criteria (i.e., hard numbers derived from field experience and engineering studies)– but that doesn't mean it's any easier.

First, there are a few things that you need to keep in mind:

- **A failure is not necessarily a hazard:** The failure of a device is only a hazard if it can cause harm to people, property or the environment through a foreseeable sequence of events. There are few medical device failures that could not conceivably result in harm – even a device that fails in a safe condition could result in a delay of treatment – but the "failure" vs. "hazard" distinction is important to keep in mind.

- **If there is no exposure to a hazard, there can be no harm[1]:** ISO 14971 states that "a hazard cannot result in harm until such time as a sequence of events or other circumstances (including normal use) lead to a hazardous situation "(Section E.1). So "exposure" is a key element in the estimation of risk – more on this later.

- **Estimating *Probability of Occurrence of Harm* requires a clearly defined harm:** In most cases, a hazard can be associated with a range of harms (e.g., a cut could lead to infection, and then illness, permanent damage to the limb, and then, possibly, death). In our previous article, *Assessing Severity*, we describe an approach for determining the harm (and associated Severity score) that is most appropriate for a particular hazard. If there are multiple, potential harms, there is likely a different Probability of Occurrence for each of those harms. Typically, harms that occur only as the result of a long sequence of events are less likely to occur than harms that are an immediate result of a device failure. You will need to clearly define the harm that is the subject of your analysis (and the sequence of events leading to that harm) to effectively estimate the Probability of Occurrence of Harm.

- **The Probability of Occurrence of Harm estimate needs to be "per use":** Too often we come across risk analyses where the Probability of Occurrence units are unstated or ambiguous. The only practical way to characterize these estimates is on a "per use" basis. For single-use devices, it's pretty simple: the probability states that the harm is likely to occur once for every 1,000; 10,000; 1,000,000, etc. uses of the device. For reusable devices, you'll need to make some assumptions about the expected life of the device (e.g., 5 years) and the number of uses over that lifetime (e.g., 250 uses per year leading to 1,250 uses over the life of the device). If your device has both reusable and disposable components, you'll need to be sure that your criteria are appropriate for both. The full benefit of making "per use" estimates becomes clear once you start collecting complaint data. With risk and complaint data described in the same terms, you'll be in a much better position to identify when gaps exist between expectations and reality.

The standard (ISO 14971:2007) brings all of the elements of risk together in one graphic in Annex E (see Figure 1). In addition to "Severity of the Harm" and "Probability of Occurrence of Harm", this graphic introduces two new terms:

- The probability of a hazardous situation occurring (P_1); and

- The probability of a hazardous situation leading to harm (P_2).

The product of those two terms is the *Probability of Occurrence of Harm*.

[1] ISO 14971:2007, Section D.3.2.2 Probability estimation.

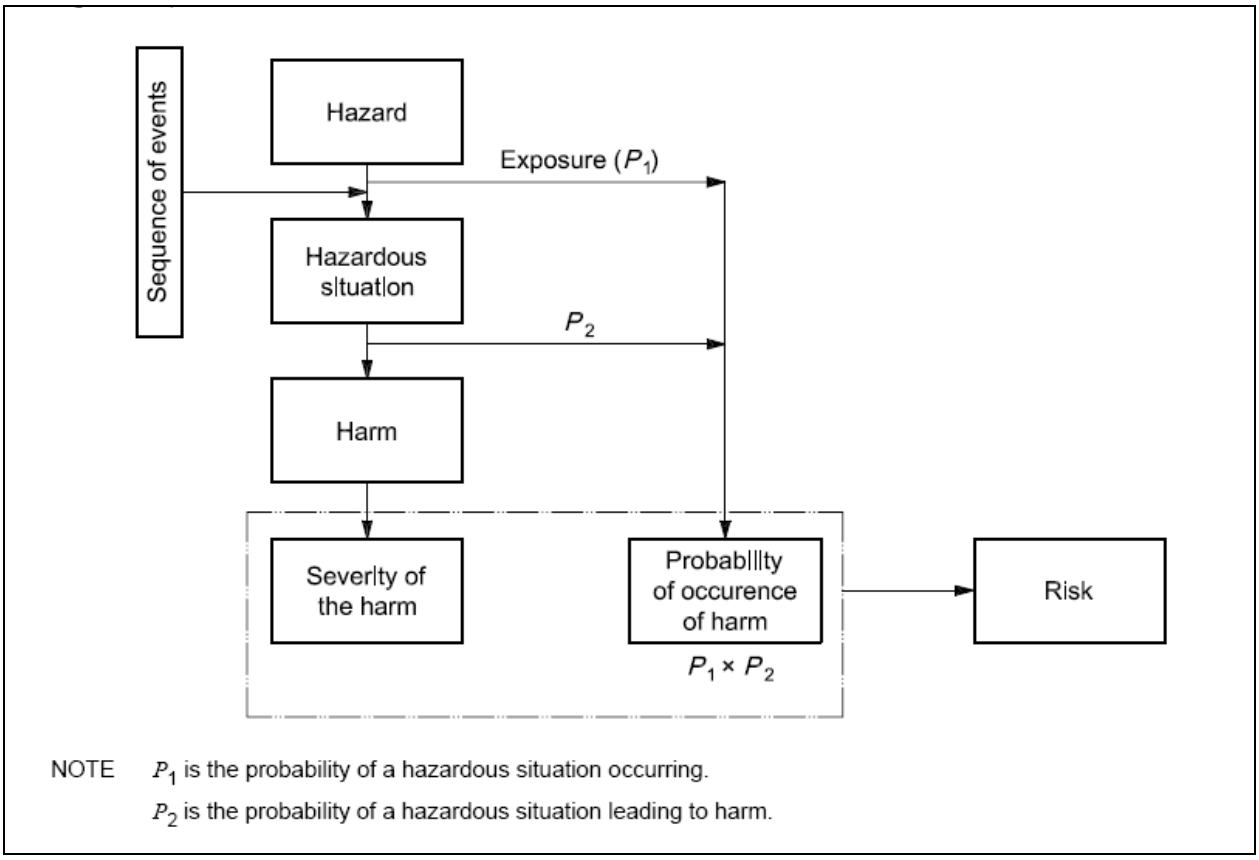

Figure 1: Components of Risk (From ISO 14971:2007 Figure E.1)

The consideration of both terms is a key element of the standard, but it is overlooked by many device manufacturers. That's not too surprising - the standard doesn't provide much guidance on how to work with these two terms.

At first, it seems pretty straight-forward . . . for example, if there's a hot surface on a device (the hazardous situation – P_1), it will only cause a burn (the harm) when the user comes into contact with that surface (the hazardous situation leading to harm – P_2). If the surface is always hot when the device is in use (P_1 = 1.0) and the user comes into contact with that surface once every ten uses (P_1 = 0.1), the Probability of Occurrence of Harm would be 0.1 - i.e., once out of every ten uses (1.0 x 0.1 = 0.1).

Most manufacturers do not explicitly estimate both P_1 and P_2 and report only a single "Probability of Occurrence" score in their FMEAs. The two terms (P_1 and P_2) are implicit in this single score. The problem is that the single score doesn't provide you with much information. Does a "Moderate" score mean that that there is a moderate chance that the hazardous situation (the hot surface) will occur and a moderate chance that it will lead to harm? Or does it mean that there a high chance of occurrence,

but only a slight chance that it will result in harm - leading to an overall "Moderate" probability of harm? There no way to tell.

Does that matter? . . . Maybe. It may affect how you decide to mitigate the risk. If the surface is always hot, a design change to reduce the temperature or place a guard over the surface may be the best approach. If the surface gets hot only once every ten uses, an alarm to warn the user may be appropriate. (Note: the alarm wouldn't make sense in the first scenario because it would be alarming all the time.)

Knowing the difference between P_1 and P_2 is also helpful when you reassess the risk following mitigation. If the mitigation reduces the probability that the hazardous situation occurs (P_1), it may have a different overall impact than if it reduces the potential for the hazardous situation to cause harm (P_2).

To better illustrate how these two factors are defined, let's look at an example based on Table E.3 in ISO 14971. Figure 2 breaks one of the five examples from Table E.3 into its specific elements. In this example, an electrostatic discharge (ESD) causes an infusion pump to fail, in turn causing harm to a diabetic patient. Four separate, potential harms are identified in the example. The one element included in our figure, but missing from the original table, is "Detection Failure." We've added that element because it seems reasonable to consider the potential for the patient to realize that he/she is not receiving treatment.

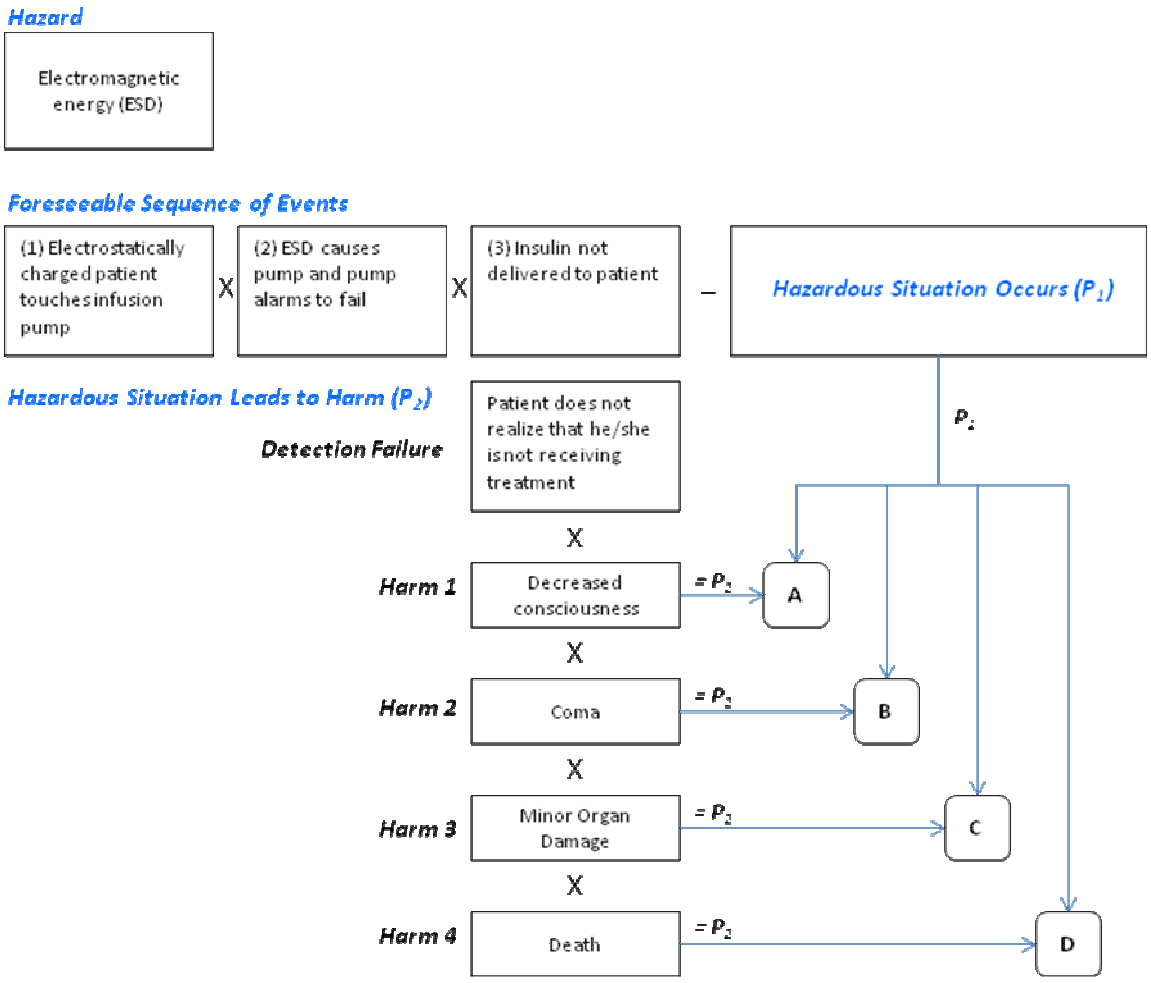

Figure 2: Probability of Occurrence of Harm Example (based on ISO 14971:2007 Table E.3)

Working from left to right, you can see that P_1, *Hazardous Situation Occurs*, is based on the probability of the occurrence of a sequence of three foreseeable events. Working down through the diagram, you can see that P_2, *Hazardous Situation Leads to Harm*, is based on the probability that the user does not detect the failure and the hazardous situation leads to a particular harm. The overall *Probability of Occurrence of Harm* is the product of the terms P_1 and P_2.

As described above and in our previous article, *Assessing Severity*, the FMEA approach to risk analysis requires that you clearly identify the harm to be considered in the analysis. As shown in Figure 2, there are four potential values for the Probability of Occurrence of Harm in this example. They are illustrated here as A, B, C or D. While P_1 is the same in each case, P_2 will vary based on the probability the hazardous situation leads to the particular harm.

Clearly, your decision regarding which harm to consider drives both the assessment of Severity and the estimate for the Probability of Occurrence of Harm. As illustrated in Figure 2, while "Decreased Consciousness" may be a relatively low severity harm, it is more likely to occur than the other potential harms. "Death", on the other hand, is a catastrophic harm, but it's less likely to occur. Is there some way to account for the full range of potential harms and associated probabilities of occurrence? Yes, but that approach creates other problems (maybe we'll get into that in a future article).

Back to P_1 and P_2 . . . While the Standard leaves the precise definitions of P_1 and P_2 open to interpretation (there are several inconsistencies in how the terms are used in Standard), the cleanest way to make practical use of these terms (while staying within the intent of the Standard) is to think of it this way:

- **P_1** is the probability of the occurrence of a sequence of events that begins with an initiating event and ends with a hazardous situation. When addressing a design or process failure, this sequence of events focuses solely on the device or process and does not consider how any outside actor (e.g., a user, operator, patient, other device, the environment, etc.) reacts to the hazardous situation. When addressing use errors, the sequence of events begins and ends with the use error that creates a hazardous situation. In all cases, there is potential for harm, but no harm has occurred.

- **P_2** is the probability that a person, property, or the environment is exposed to the hazardous situation and is harmed as a result. In some cases, the sequence of events is pretty short and clear, e.g., following the failure of a life sustaining medical device - the patient dies (P_2 = 1.0). In other cases, the sequence of events may be complex, e.g., following the failure of an infusion pump and pump alarms due to ESD (insulin is not delivered to patient when needed) - the patient does not realize that treatment is not being delivered, he or she suffers from decreased consciousness and does not receive immediate treatment, and then suffers from a series of harms including coma, and ultimately death (P_2 = something less than 1.0).

What's helpful about this approach is that it separates the engineering-based estimate of the probability that the hazardous situation occurs (P_1) from the sometimes complex series of events associated with the occurrence of harm (P_2). Often, P1 can be based on Reliability Engineering estimates (e.g., the mean time between failure of the valve is 100,000 uses, P_1 = 0.00001). These probabilities are relatively objective engineering estimates based on test results and design standards.

On the other hand, estimating P_2 requires an understanding of the use environment, user profiles, and the likelihood that the hazardous situation will lead to a particular harm. Included in this sequence of events are any opportunities for the hazardous situation to be detected before the harm occurs. From this perspective, detection refers to any point in the process (following the creation of the hazardous situation) where the problem can be identified and the harm is prevented. There may be multiple opportunities to detect the hazardous situation. For example, Annex H of the Standard describes the

scenario where an incorrect IVD result (the hazardous situation) causes harm only when it fails to be 1) detected by laboratory staff and 2) detected by the healthcare provider.

So how does this work in practice? In Table 1, we've provided an example that builds upon the ESD/Infusion Pump example from above. Our fictional company's engineers have come back to us with an estimate for P_1 of 1×10^{-6} (i.e., once in one million uses), and our risk analysis team estimates that P_2 is 5×10^{-5} (i.e., five times for every 10,000 hazardous situations, or once in 2,000 situations).

Using these contrived probability estimates, the Probability of Occurrence of Harm is 5×10^{-11} (i.e., 5 times for every 100 billion uses, or once for every 20 billion uses).

Table 1: Probability of Occurrence of Harm Calculation (based on ISO 14971, Table E.3)

P1 Calculation		**P2 Calculation**	
Event	**Probability of Occurrence**	**Event**	**Probability of Occurrence**
1) Electrostatically charged patient touches infusion pump	0.01 [once in 100 uses]	User does not detect that no treatment is being delivered	0.5 [half the time]
2) ESD causes pump and pump alarms to fail	0.0001 [once for every 10,000 times that an electrostatically charged patient touches an infusion pump]	Decreased consciousness	1.0 [every time that treatment is not delivered]
3) Insulin not delivered to patient	1.0 [every time that the pump and alarms fail]	Minor organ damage and/or coma	0.01 [once for every 100 times that treatment is not delivered]
-	-	Death	0.01 [once for every 100 patients with minor organ damage and/or coma]
Total: Probability that Hazardous Situation Occurs (P_1)	$1 \times 10^{-6} = 0.01 \times 0.0001 \times 1.0$ [once in one million uses]	**Total:** Probability that Hazardous Situation Lead to Harm (P_2)	$5 \times 10^{-5} = 0.5 \times 1.0 \times 0.01 \times 0.01$ [5 times in every 10,000 hazardous situations, or once for every 2,000 hazardous situations]

Note: All probabilities are fictitious and intended only to illustrate the analysis approach.

One potential death in 20 billion uses sounds like a pretty low number. But if the device delivers insulin five times a day, 365 days a year to 5 million users (i.e., about 9 billion uses per year), there's about a 50% chance that a user will die from this specific failure each year. Is this acceptable? That depends on your company's perspectives on the risks and benefits of the device - as established in you risk acceptability criteria. *We'll address that part of the process in a future article*.

To help simplify the analysis, the Standard recommends establishing probability levels (typically 3 to 5 levels) so that harms that occur at a similar rate can be grouped together. Table 2 is the example table provided in the Standard.

Table 2: Example of Semi-Quantitative Probability Levels (from ISO 14971, Table D.4)

Common Terms	Examples of Probability Range	Layman's terms
Frequent	$\geq 10^{-3}$	More than 1 per 1,000
Probable	$\leq 10^{-3}$ and $\geq 10^{-4}$	Between 1 per1,000 and 1 per 10,000
Occasional	$\leq 10^{-4}$ and $\geq 10^{-5}$	Between 1 per10,000 and 1 per 100,000
Remote	$\leq 10^{-5}$ and $\geq 10^{-6}$	Between 1 per 100,000 and one in a million
Improbable	$< 10^{-6}$	Less than one in a million

So where would the Probability of Occurrence of Harm from our example land in this table? It would be many magnitudes below the table's "Improbable" category. Does this mean that the probability of harm in our example is improbable? Not at all.

We already identified that a patient death is likely to occur once every other year. As the Standard points out, "the definitions for probability can be different for different product families. For example, a manufacturer can choose to use one set of definitions for X-ray machines, but can have a different set of definitions for sterile disposable dressings." It then goes on to identify several questions that you will need to consider when establishing your probability ranges including: "How often is a particular medical device used?" and, "What is the number of users/patients?"

Given the high volume and continual use of the device in our example, it would be important to establish probability ranges that effectively capture the levels of occurrence that we are likely to see once the product is on the market.

This seems like a lot of work . . .

Yes, it can be. The approach described in this article sits somewhere between assigning a single Probability of Occurrence of Harm score and completing a full Fault Tree Analysis - see our earlier articles (numbers 4 and 5) that describe the use of FTA.

The question is what level of rigor is needed to produce an effective risk analysis – and by "effective analysis" we mean an analysis that helps you make effective decisions for how to improve the safety of your device. As we described in an earlier article on setting priorities (Risk Management #2), before your team jumps into the details of an FMEA, you should take a high-level view of the risks associated with your device to identify the high priority areas. The rigor that you apply to the analysis of specific aspects of your device and its uses should be proportional to the significance of the potential risks.

If the only imaginable harms associated with your device are of "Minor" severity, you may be able to go right to your FMEA and not have to worry too much about dissecting the Probabilities of Occurrence of Harm into all of its component elements. But if there's any indication that your device could cause a serious injury or death, it's your responsibility to identify those priority areas and complete a thorough analysis. That analysis may involve thinking through P_1 and P_2 as we've discussed in this article, constructing FTAs to better understand how best to mitigate particularly complex hazardous situations, or applying some other tool that is appropriate to the particular situation. The key is that you use these tools to gain as full of an understanding of the device risks as possible before releasing your product to the market -- so that you can provide patients and users with the safest possible products.

Next Steps

This article and the Risk Management #5, *Assessing Severity*, addressed the key elements of risk analysis. More could be said about both, and we may go back at some point and dig a little deeper. For now, we'll move on to the next step in the process and address how these elements of risk can best be used to improve product safety.

If you've missed any of the previous articles look them up on www.medicept.com/blog/.