Probabilistic risk analysis, Monte Carlo simulations, and Federal Communications Commission decision making

DLA PIPER

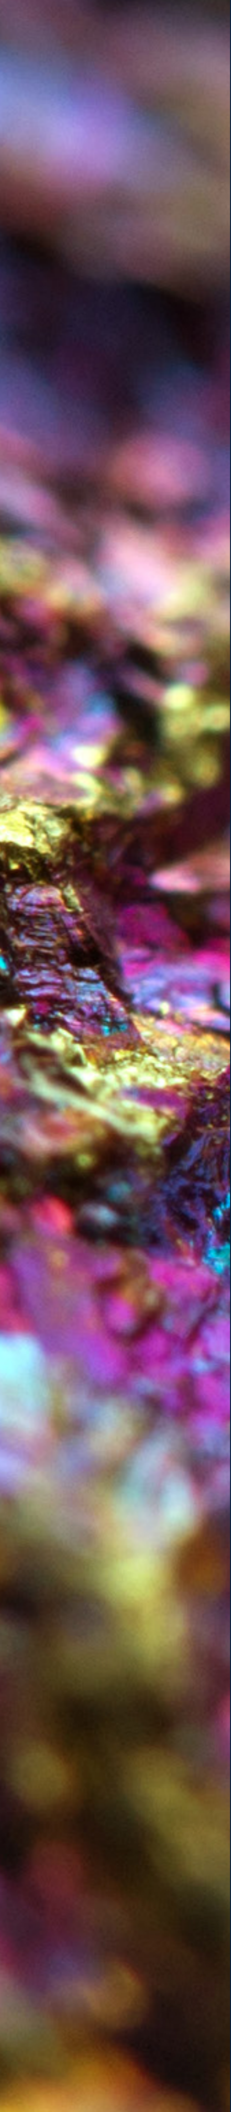

Probabilistic risk analysis, Monte Carlo simulations, and Federal Communications Commission decision making

In recent years, probabilistic risk analysis (PRA) – which includes techniques like Monte Carlo simulations – has become a recurrent theme in various public and private spectrum policy initiatives involving the Federal Communications Commission (FCC or Commission) and the National Telecommunications and Information Administration (NTIA). PRA techniques promise more informed decision making in complex areas like inter-service radio interference assessments.

This article discusses what PRA is, how Monte Carlo simulations work and fit within the PRA, statistical concepts relevant to how PRA is used in practice, and how PRA could alter radio policy decision making.

What is PRA?

PRA is typically contrasted with decision making techniques that rely on simple worst-case scenarios. Instead of answering regulatory questions with a binary "yes" or "no," PRA seeks to characterize situations in a more nuanced manner that considers the chance that certain events take place.

An example of worst-case decision making might be declining to schedule a picnic on an upcoming Saturday because it *could* rain. PRA-informed decision making instead might evaluate whether to schedule a picnic based on the probabilities associated with rain. In this instance, the analysis could be very simplistic, *eg*, "there is a 50 percent chance of rain between noon and 2 pm on Saturday."

The analysis could also be much more complex, considering not only the risk of rain, but also the severity and duration of the rain, whether nearby shelter is available, or whether the picnickers are indifferent to rain. A characteristic of PRA is thus that it does not lead to simple regulatory outcomes, but instead allows a weighing of the costs and the benefits of a particular regulatory decision.

PRA is not new. Federal organizations have explicitly used PRA for many years regulating various activities, many with potentially perilous or costly outcomes. For example, the Environmental Protection Agency (EPA) has defined guidelines for PRA – as discussed further below – and utilizes the techniques when considering, among other things, health effects from human exposure to toxic materials. The Nuclear Regulatory Commission uses PRA when "estimat[ing] the frequency of accidents that cause damage to the nuclear reactor core" for US powerplants, and the National Aeronautics and Space Administration utilizes PRA, among other things, "for crew transportation system missions to the International Space Station (ISS)."

Arguably, the FCC has always applied some implicit PRA in its policy decisions – by acknowledging a negative static outcome documented by a commenter, but ultimately dismissing the risk of the outcome as unlikely. What would be new – at least for the FCC – is adoption of recognized best practices for how PRA is used.

What is a Monte Carlo simulation?

A Monte Carlo simulation is a mathematical algorithm that determines the likelihood of outcomes using repeated random sampling. Monte Carlo simulations are often linked with PRA, although PRA is a broader concept. In an ideal world, a regulator might have a formula that precisely describes the relevant risks associated with a particular decision. To the extent such a formula exists, taking actions based on the described probabilities would be PRA-informed.

Unfortunately, most complex regulatory decisions involve scenarios with a large collection of probability distributions that make derivation of precise risk formulas impractical. In such cases, Monte Carlo simulations can help a regulator to understand the risk probabilities by employing a statistically valid random sampling of variables and determining specific outcomes in those cases. Thus, Monte Carlo simulations are often associated with PRA.

To illustrate how a Monte Carlo simulation works, consider the possible values from rolling a pair of dice, values that range from 2 to 12. If the probability distribution is defined with only a half dozen rolls, shown in Figure 1, the simulation is not particularly precise.

Figure 1: Monte Carlo with 2 Dice and 6 Rolls

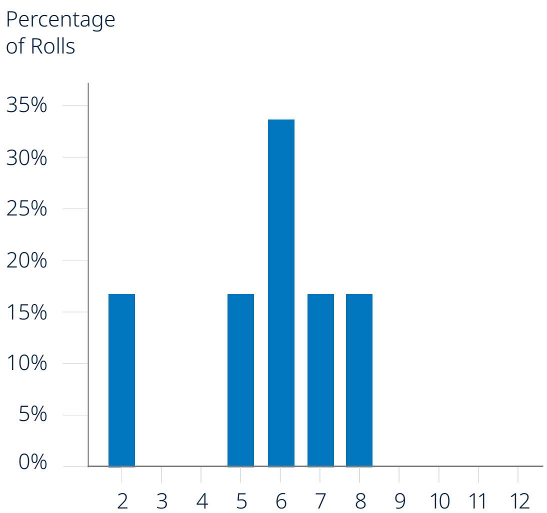

On the other hand, if the simulation is run looking at 50 rolls, the resulting probabilities start to become more representative:

Figure 2: Monte Carlo with 2 Dice and 50 Rolls

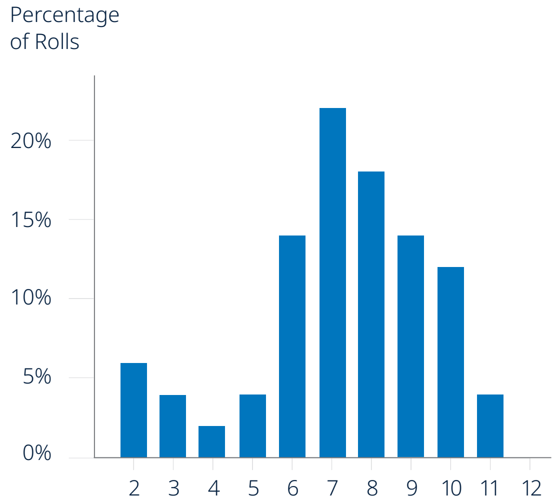

And with a very large number of rolls, the simulation becomes quite accurate:

Figure 3: Monte Carlo with 2 Dice and 100,000 Rolls

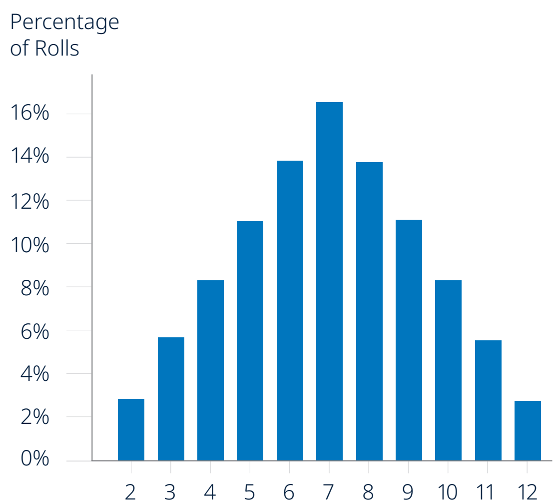

In this particular instance, the two probability distributions are relatively easy to describe as mathematical constructs, and there are few variables. As a result, the distribution could be described with a formula that is absolutely accurate. But when the number of distributions that are inputs to the model becomes much larger, ie, the simulation considers more factors relevant to a particular analysis, then converting the collection of individual distributions into a single distribution becomes quite complex. In those scenarios, it is often easier to leverage computing power with a large sample set to understand the distribution with a sufficient degree of statistical certainty.

An aside on statistical concepts for PRA and the benefits of PRA

To fully appreciate the impact of PRA, it is useful to understand a few basic statistical concepts – in particular, the "normal" distribution and cumulative distribution functions (CDFs). Each of these is discussed briefly below.

THE NORMAL DISTRIBUTION.

The dice-rolling in the prior discussion is an example of a "discrete" distribution – each number on a die has an equal probability of occurring. In the real world, probability distributions are less likely to be discrete, but rather vary from some median or average value – the so-called "normal" or "gaussian" distribution.

Perhaps the most widely experienced normal distribution is bell curve grading. In strictly numerical, non-bell curve grading, letter grades are assigned based on fixed numerical ranges on an exam – every score of 90 percent or better is an A, every score from 80 percent to 89 percent is a B, and so on. Grading on the curve, in contrast, takes the average score in the class and assigns letter grades based on deviations from that average. Assuming a class of "n" students, when the traditional bell curve is applied, the professor takes the average grade in the class (the mean, represented as μ, and equal to a score of 40 in Figure 4), and every student within one standard deviation (represented as σ and, in this example, equal to 7) of the average is awarded a C, every student that is above the average by one to two standard deviations gets a "B, and so on, as shown in Figure 4 below:

Figure 4: Traditional "Bell Curve" Grading

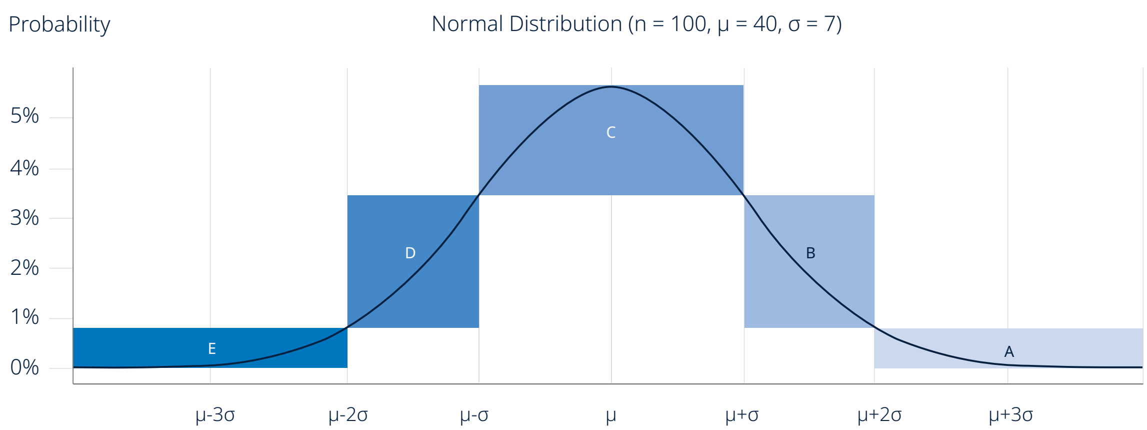

In other words, the students with scores in the range from 33 to 47 (μ ± σ, or 40 ± 7) would get Cs, those with scores in the range from 47-54 would get Bs, those with scores of 54 and above would get As, those with scores of 26-33 would get Ds, and those with scores below 26 would get Es. The standard deviation is defined so that, for a bell curve, approximately 68 percent of the values lie within one standard deviation of the mean (*ie*, ±σ, meaning that approximately 68 percent of students get Cs), approximately 95 percent lie within two standard deviations (±2σ), and nearly all are within three standard deviations (±3σ). Normal distributions are, as noted, more common in the real world, and variations of normal distributions might be used to model things like transmitter power, antenna height, channel size, or other factors relevant to interference assessments.

CUMULATIVE DISTRIBUTION FUNCTIONS.

In PRA, it is generally less useful to know what percentage of instances a simulation arrives at a particular value – rather, the important figure is the percentage of instances the value is *at or below* a particular value.

For example, if the FCC is evaluating a scenario in which the signal strength threshold for harmful interference is 100, the useful metric is the percentage of time the signal strength is below 100, and not, for example, the percentage of time the signal strength is specifically at, say, 78 or 99. This is where the cumulative distribution function or CDF comes in. The CDF is essentially a probability curve that shows, on the y-axis, the percentage of time a value is at or below the value on the x-axis – the CDF cumulates, or adds up, the probabilities for any value lower than a particular value on the x-axis. As a result, the right-most and top-most part of the CDF should be 100 percent, designating that all values are at or below the highest possible value.

The dark blue line in Figure 5 shows the CDF associated with the class grading distribution detailed in Figure 4 (which is reproduced in Figure 5 as the bright blue line):

Figure 5: CDF versus PDF

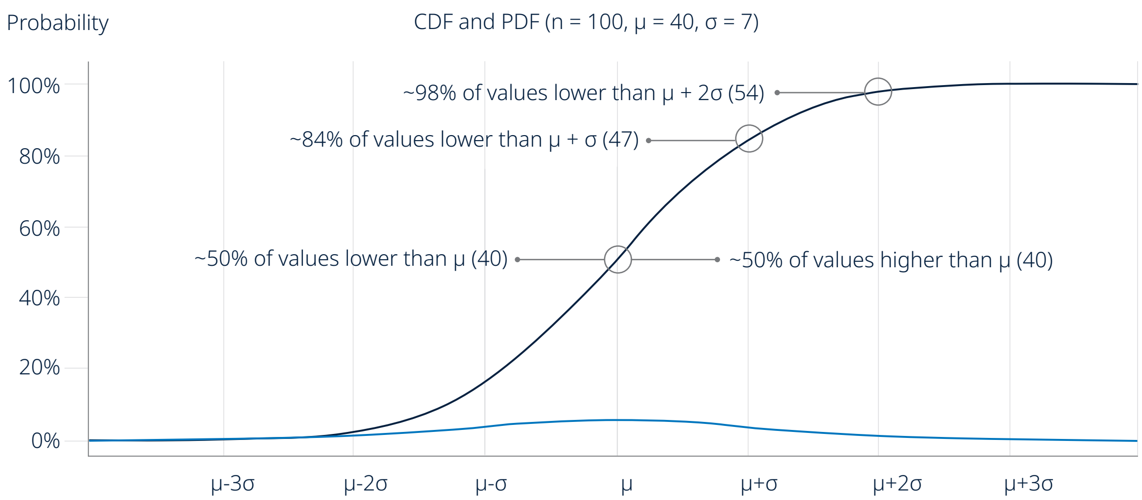

Even though the bright blue line is the same probability distribution function (PDF) shown in Figure 4, it looks different because the vertical scale has changed: whereas the top of the vertical scale in Figure 4 is approximately 6 percent, the top of the vertical scale in Figure 5 reaches 100 percent. As you would expect with a normal distribution, and as shown with the grey circles, the dark blue CDF line shows that 50 percent of the grades are below the average (μ) and 50 percent above the average. The distribution also shows that approximately 84 percent of the grades are Cs or lower (below μ + σ) and approximately 98 percent are Bs or lower (below μ + 2σ).

How can PRA better inform decision making? The CDF offers regulators more than a simple yes or no, and the shape of the CDF that can potentially offer more informed regulatory decisions. Consider the dark blue and bright blue CDFs shown below in Figure 6.

Figure 6: Comparison of two CDFs

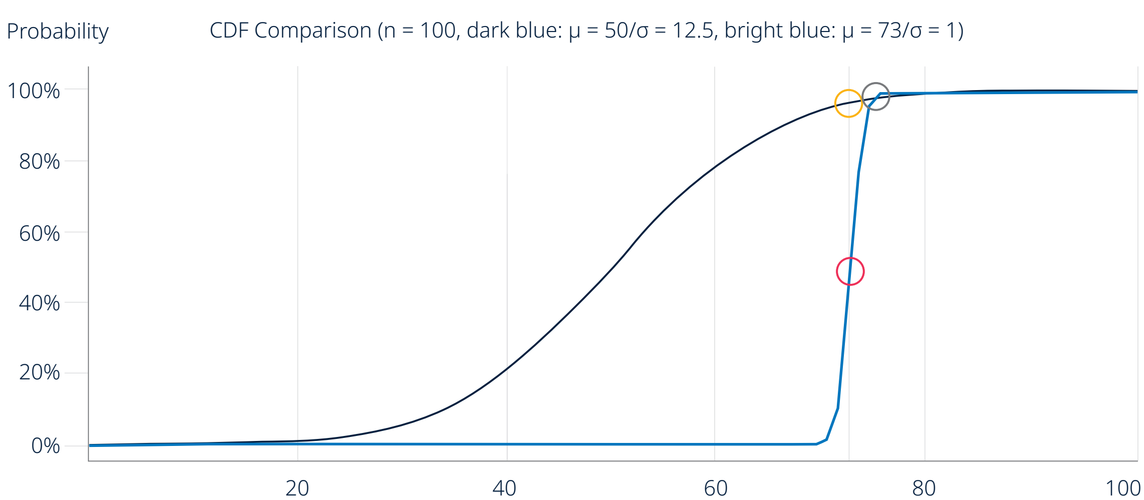

Both are normal distributions, so the point designated by the grey circle shows in both cases that, roughly 98 percent of the time, the value is less than 75. Specifically, the "$\mu + 2\sigma$" point (which is correlated to a probability of ~98 percent) for the dark blue curve is $50 + (2 \times 12.5)$, which is 75, and the "$\mu + 2\sigma$" point for the bright blue curve is $73 + (2 \times 1)$, which also happens to be 75.

Figure 6 illustrates that the consequence of an error or miscalculation in the simulation is far more dire in the bright blue case. If the assumed interference threshold was 75, the Commission could conclude that both simulations show interference is unlikely – a roughly 2 percent chance (as shown by the point designated by the grey circle). If it turned out the *actual* threshold for harmful interference was 73 instead of 75, however, the point designated by the yellow circle shows that the change in the probability of interference for the dark blue curve might still be considered relatively benign – approximately 95 percent of the cases would still be below the interference threshold. But for the bright blue curve, as shown by the point in the red circle, the percentage of scenarios where harmful interference occurs would rise to 50 percent. This means basing a decision on the bright blue curve carries greater inherent risk than in the dark blue case and illustrates why PRA can potentially lead to more informed policy decisions.

Another benefit of PRA is that, because PRA simulations are typically computational models, the inputs can be varied and the resultant changes to risk analyzed. For example, as shown in Figure 7, if one of the criteria in a simulation is modified and the CDF changes from the dark blue line to the bright blue line, a regulator might conclude that varying that criterion entails little added risk:

Figure 7: Evaluating changed assumptions

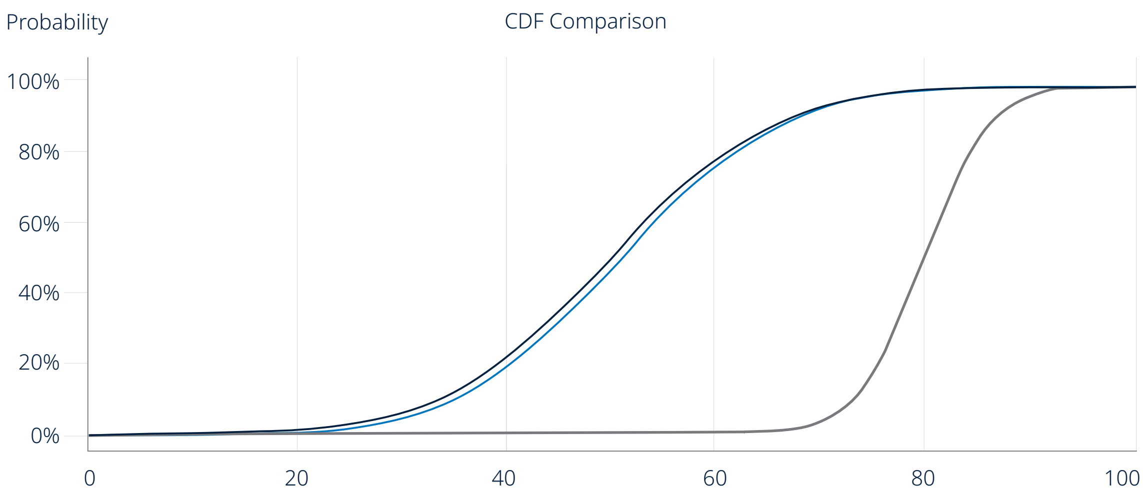

On the other hand, if varying the criterion results in the more radical change from the baseline dark blue CDF to the grey CDF, the regulator may conclude that adopting a rule to regulate that criterion is warranted.

A final advantage to the use of PRA is that the techniques permit the modeling of extremely complex scenarios – and the Commission is increasingly being asked to consider prospective spectrum use scenarios that involve exponentially more convoluted and interrelated criteria. Not only have radio technology and techniques evolved to make more efficient and effective use of limited bandwidth resources (which affects both the transmitter and the victim receiver), new propagation models consider a wider range of physical path characteristics, the data on those characteristics (topology, land cover, building and obstruction data) is more accurate and at a higher degree of resolution, spectrum access mechanisms provide greater capability to engage in real-time coordination, and, perhaps most importantly, the need for spectrum is driving consideration of spectrum sharing scenarios that create potential radio interactions that are more intricate. In this environment, the Commission will inevitably need tools that better illuminate the costs and benefits of new radio uses.

Formalizing PRA in the context of Commission radio interference decisions

Historically, the Commission's assessment of potential harmful interference between a proposed new radio use and existing incumbent use was often modeled as a static, worst-case scenario. In other words, a hypothetical would be defined that involved an incumbent user, a new user, and geometry and link characteristics that would create the worst possible, but realistic, environment for the incumbent. Although scenarios like that can be valuable in understanding possible impacts of new allocations, the knowledge that harmful interference is *possible* is generally less useful as a regulatory tool than understanding the *risk* of interference. But PRA will only improve decision making if it is employed in a principled manner.

There is considerable learning on how to apply PRA in a principled manner. Underscoring that PRA is neither new nor radical, 25 years ago the EPA approved a document titled "Guiding Principles for Monte Carlo Analysis" that was "part of a continuing effort to develop guidance covering the use of probabilistic techniques in [EPA] risk assessments." The document outlines key considerations when Monte Carlo simulations are used in rulemakings—considerations that can easily be adapted to Commission rulemakings. These guidelines include, among other things:

1. **The purpose and scope of the assessment should be clearly articulated.** If a probabilistic analysis is intended to illuminate the risk of interference, the analysis should obviously define the problem that is being evaluated. Specifically, the study needs to identify the class of devices that are the interferers and the victims and any assumptions about the scenario being testing – eg, whether the study assesses cumulative interference or single source, whether the study incorporates busy hour considerations or average use, how the study defines "harmful interference" for purposes of testing, and whether there might be other contributors to the interference threshold that are not considered. As an example, if the scores in Figure 5 were grades that were intended to reflect the objective level of knowledge of a class, the CDF would have very different real-world meaning if the class in question was composed entirely of elementary school children, as opposed to being composed entirely of doctorates in that subject.
2. **The analysis should be independently reproducible.** In the EPA's words, "the methods used for the analysis (including all models used, all data upon which the assessment is based, and all assumptions that have a significant impact upon the results) are to be documented and easily located in the report," including "the names of the models and software used to generate the analysis." In an ideal world, PRA and Monte Carlo analyses should be evolutionary – a proponent of a new service or a party requesting a waiver may model a scenario as best it is able, but that model could well be improved incrementally throughout a rulemaking as new factors relevant to the analysis are raised and incorporated into the model.
3. **The analysis should incorporate a sensitivity analysis.** Specifically, "probabilistic techniques should be applied to the . . . factors of importance to the assessment, as determined by sensitivity analyses or other basic requirements of the assessment." As previously mentioned, the sensitivity analysis can be a mechanism that allows regulators to better regulate—for example, a regulator might better balance spectral efficiency and noninterference by allowing more or less flexibility with respect to technical regulations depending upon the actual relationship of those factors to modeled harmful interference. If an analysis suggests that harmful interference in a particular situation is less driven by power and more driven by antenna directionality, that knowledge can allow more targeted and precise regulation. In simplistic terms, a rigorous probabilistic simulation offers regulators the ability to adjust a number of regulatory dials or levers—adjustments that ultimately might form the basis of regulations – and see the effect those changes have on the outcome.
4. **The analysis should identify and account for correlations and dependencies.** When building a complex statistical model, understanding correlation and dependency can significantly alter the results. For example, if one studied traffic at an intersection and found that one in four cars are red, one in four cars are sports cars, and one in ten cars speeds, a modeler might create a simulation where there is a 1 in 160 chance a car passing through the intersection is: (i) red; (ii) a sports car; and

(iii) speeding. But if it turns out all sports cars must be red – a dependency – the chance should drop to 1 in 40. And if it turns out that sports cars are more likely to speed than other cars—a correlation—the probability would drop even further.

5. **The input and output assumptions should be documented.** When specific inputs are used in a model, the assumptions underlying how those inputs were modeled should be not only documented, but "explained and justified," including discussion of "variability and uncertainty." The adage "garbage in, garbage out" is absolutely valid in the world of statistical modeling, so understanding what is being fed into a model is a crucial component in assessing the validity of the model. When the outcome of two dice being thrown is modeled, the input assumptions are straightforward – each die can be an integer value between one and six, and each value has an equal probability of occurring. But real-world models are rarely that simple – if a key factor in an analysis is antenna height, basing antenna height on a database of existing height values may seem reasonable. But it may not be reasonable if, for example, it turned out that the historical database had only three entries, or that the service was undergoing a change to a distributed architecture where tall antennas were being rapidly phased out in favor of lower antenna heights.

6. **The statistical reliability of the study should be assessed.** In the EPA's words, "[t]he numerical stability of the central tendency and the higher end should be presented." In other words, if the results are subject to significant change with new iterations, such as the rolling the dice a seventh time and comparing that to the distribution seen for 6 rolls in Figure 1, it is an indication that not enough sampling has been done. In such respects, the "tail" is important because the tail of the distribution is often where the negative consequences occur – eg, interference exceeding some high threshold only occurs in a very small number of cases. The relevance of the tail is that the "tail" cases may warrant some examination to see the mechanism of interference, so if there are not a number of cases where the tail conditions are satisfied, a full understanding of the interference effects may be lacking. It may sound good, for example, that interference occurs only in 0.1 percent of cases. But if that 0.1 percent represents one iteration out of one thousand evaluated, it may not provide the same degree of insight into interference as the situation where there are 1,000 iterations showing interference out of 1,000,000.

7. **Probabilistic analysis should consider deterministic results if possible.** As it is the goal of PRA to evaluate the real risk posed, PRA should be validated using actual measured data to "[p]rovid[e] . . . comparisons between the probabilistic analysis and past or screening level risk assessments." The EPA also recognized that "deterministic estimates may be used to answer scenario specific questions and to facilitate risk communication." Thus, PRA should remain grounded in real-world performance—to the extent that the real-world performance is at variance with the results predicted in PRA, those inconsistencies should be explained, or the model adjusted.

The EPA's guidelines, and similar work by other federal agencies, also contain a wealth of other practical information on PRA methods, including some discussion on when PRA is useful and worth the expense. When complex new radio services are being introduced, especially in a mobile environment, PRA techniques may well be justified and could lead to more efficient utilization of the radio spectrum.

In conclusion

We hope this provides some useful understanding of the core concepts in probabilistic decision making, and how those techniques might be employed at the FCC. If you would like to know more or would like to build a Monte Carlo simulation to better understand the risks associated with a regulatory activity, do not hesitate to contact the Telecoms Engineering team at DLA Piper or either of the authors:

Eric DeSilva
Partner
eric.desilva@dlapiper.com

Emma Marion
Associate
emma.marion@dlapiper.com