

The Role of Risk Management in Regulation

September 5, 2018

We live in a society that assumes risks can be both anticipated and controlled. If there is a failure to foresee and prevent harm (to fail to control the risk) and a regulator in sight, odds are it will be on the list of those to blame. And the standard of affixing blame to a regulator is low. Facing these pressures, regulators are using management tools and concepts to inform their operations and programs. This article is a quick tour through the role of risk management in a regulatory context.

Risk management as a formal method dates from after WWII. It began in the world of finance and engineering and involved (and still does) technical and quantitative models and approaches. At some point these formal methods were adopted for use outside of financial and technical contexts. At the same time, private firms were beginning to apply risk management to the entire span of their operations. In addition to the risk of financial loss from such things as market or interest rate changes, firms started to look at other sources of loss such as legal or political hazards. This was the concept of enterprise risk management or 'ERM' for short.

If you are a regulator looking for a risk management program, what experts will likely try to sell you on is what I would term, "ERM lite". "Lite" refers to the stripping away of the quantitative analysis that goes along with much of the risk management work in finance, insurance and engineering. A barebones ERM process would look something like this: First, you set the context. This involves setting the scope of the ERM process; 'scope' being what areas and objectives the process will apply to. ERM tries to cover most facets of the organization: finance, operations, legal and so forth. Context also includes setting the risk tolerance of the organization in each of these areas. Understanding how the organization and external groups have reacted to harmful events in the past is a key step in setting risk tolerance.

The next step is to identify risks, or events that may have an adverse effect on the organization or its objectives. There are many different ways to do this and organizations usually use more than one method. These include workshops, surveys or more detailed methods like scenario review or case studies. The process usually considers each area of risk separately. For example, what can happen from a legal perspective to affect this or that objective, or how would a certain event affect the finances of the organization. The outcome is a list of risks (negative events) organized across the different risk areas.

With the risks on the table, the next stage is assessment. The goal is to rank the risks to determine priority. Assessment is about two things: the likelihood that the event will happen and how severe is the outcome of the event. Likelihood and severity both receive a score using a similar scale. The scale can be a numerical rank, for example 1 through 5. A scale may instead use descriptive words, for example low, medium and high. The risk level is then calculated from the two scores. By way of example, a risk with high likelihood and high severity ranks higher than a risk with low likelihood and medium severity. It is easy to map the risk assessment using a matrix, with likelihood on one axis and severity of impact on the other.

Things get a little more complex than this because the assessment should consider what is already in place to control the risk. What is left over after the existing controls is the 'residual risk', which should drive the next step: risk treatment.

What you have now is a list of risks ranked according to their combined score of severity and likelihood. Risk treatment is the process

of choosing ways to respond to the risks. The organization may choose to accept risks below a certain score while coming up with controls for the highest risks. There are various ways to treat or respond to risks. Examples are to transfer a risk to another party, to avoid or reduce it or to insure against it. The risk tolerance of the organization will determine the type and extent of the treatment.

Once assessment is complete, the risks and responses to them are set out in a risk register. This provides a basis for tracking the risks. The risk scores and treatment will require adjusting from time to time. Another important feature of the process is communication about risks and responses to them, both inside and outside the organization.

That in a nutshell is ERM. It certainly seems at first glance to be a good thing. An organization should know the risks it faces and try to control them. But it is hard to find evidence that it works or that it actually reduces risk. And it has a high cost to prepare and implement. There are also a number of examples of things going very wrong with a lot of risk management in place. The Deepwater Horizon blow-out or the Space Shuttle Challenger disaster, for instance. This suggests that a risk management process alone will not suffice (and you also need to build the right culture, which of course is more time and effort).

These issues aside, it is a good idea for an organization to know its major vulnerabilities and have plans in place to avoid them or deal with the consequences. Given the time and cost involved it is hard to know how much to invest in a formal ERM process. But this is more of a question of governance than regulation. What I want to address here is the potential use of ERM methods to inform the regulatory program.

If you are an organization with a regulatory mandate a vital question to ask is, "risk to whom?" The focus of ERM is on risks to the organization. Its purpose has largely been to identify and control threats to the bottom line or to achieving certain goals, or to ensure a firm can stay afloat during a crisis and maintain public trust. But regulators have a mandate to protect the public. How much sense does it make to use a process that is more about preserving the organization than protecting the public? In response, one could simply say that the role of ERM could be to ensure a regulator achieves its mandate to protect the public. In this context, a regulator would use ERM to identify and control risks to the public.

This seems to make sense. But given the cost of implementing ERM it may be worthwhile to look at some of its shortcomings. Not all of these shortcomings relate to cost. Some of them go to the issue of whether the simplistic approach to risk assessment in ERM is a sufficient basis for a regulatory program aimed at protecting the public. Here are five things to think about before embarking on a risk management program:

1. It takes a lot of resources

The first consideration does relate to cost. Risk management can turn into a large undertaking. It can be a heavy user of staff time. Because it calls for a high level of analysis and often technical skills you may need to recruit or retain more talent, including an individual with expertise in analyzing data. It also tends to be report heavy and by its nature needs a high degree of monitoring. A big risk of risk management turns out to be its opportunity cost. The (hard) decision is often about what existing projects to drop in order to do risk management properly. My anecdotal observation is that risk management projects often get off to a good start. But sustaining and completing them is a larger challenge.

2. Data comes first

You need data about a sector to identify risks to the public. Most importantly, you need evidence relating to what harms the public face in a certain sector. This data has to be reliable and accurate. It should also provide a general enough picture of the distribution of harms within the sector. It is true that there are ways to identify risks without the help of statistical data. This would include relying on anecdotal or comparative evidence. But it is hard to think that risk management would be of much value without good sources of empirical data. This point concerns the platform or starting point for risk management. If you don't have data, you should likely create

the means of collecting and analyzing it before investing much in risk management. And this is by no means easy if you are a complaint-based regulator. You only see largely what the public brings you and you must be creative in searching out and finding harms that do not surface through complaints.

3. Agreeing on riskiness

Let's assume you are fine with the cost and you have data. You will then have to find some sort of relatively objective standard for determining risk. People find it hard to agree on two things: the degree of risk (of something or somebody) and what level of risk to tolerate. The first relates to what is or is not risky. This is notoriously subjective and it can take a long time to agree upon. The second relates to what amount of risk is acceptable. This involves trade-offs between cost and safety. A risk management program will not tell you how much you should invest in public safety. But regulators have to answer this when deciding which risks to respond to and how much to invest in response and control for any given risk. They also have to decide how much time and resources to spend on identifying risks. In other words, they must decide when to stop identifying risks and when to move on with the job of controlling them. These challenges of risk perception and striking a balance between cost and safety are not reasons to reject risk assessment. But they are an added level of complexity and take time and effort (and cost) to sort through.

4. Risk assessment may be too simplistic

The assessment stage in a risk management program looks at two things: the likelihood and the severity of a risk. Of course these are important factors to consider. But risk management tends to be a 'top heavy' or 'early stage heavy' approach. Risk identification and assessment result in a long list of risks and then the treatment stage simply says "control them". But a successful control strategy will require much more than an estimate of how likely the event or harm may be and how bad it will be. To intervene and prevent harm the regulator will need to understand the harm in some detail. What causes it? What kind of response or treatment will work? These are implementation and practice questions. And risk management programs spend very little time down in the weeds, so to speak, with respect to these details. The danger is sinking too much effort into developing a list of risks with little resources left to understand them in depth, let alone trying and testing possible solutions. Perhaps regulators would do better to identify and tackle a short list of well understood problems rather than compiling a long list of more cursory risks.

5. The perils of prediction

To assess a risk is to forecast it. Both the likelihood and severity of a risk will by and large involve an estimate. But history and science tell us that most people are terrible at forecasts. There is no shortage of biases that hamper prediction. Groupthink, risk anchoring and framing, and the recency effect are all examples of pitfalls that stand in the way of accurate forecasts and estimates. And added to these pitfalls are Black Swan (remote chance, high impact) events that are impossible to predict. These will occur no matter how sophisticated the risk management program. The failure to confront the problems of prediction is the Achilles Heel of most risk management programs.

It is worth spending some more time on the role of prediction in the regulatory context. More and more regulators have started to use risk assessment as a tool in their programs. An example is the use of risk scores to guide discipline decisions or outcomes and to target inspections. This is different than trying to predict the risk of harm. This is about predicting who is at greater or lesser risk to harm the public. And this may be an area where theory and method is moving faster than the evidence of effectiveness and the safeguards in place for fairness and due process.

The field of law enforcement is a great place to look to learn about the good and bad of person-based targeting. Police are using big data and other new technologies to predict and deter crime. One approach is police 'heat lists'. These lists rely on multiple factors that combine to create risk scores. The factors include criminal history, arrests, parole status and gang membership. A high score means a greater chance of being a victim or committing a violent crime. In Chicago, being on the heat list resulted in a visit and written notice

from the police. The intent was to deter future crime by warning the people on the list that they are under watch and that unless they avoid future crime they will face the full force of the law.

The Chicago heat list has had at best mixed results. A third-party review by the RAND Corporation found that it had little predictive accuracy and effect on violent crime. Problems with prediction are not limited to the criminal context or to targeting individual people. In the UK, the Care Quality Commission (CQC) developed a statistical tool to rank hospitals according to risk and to decide which ones to inspect first or more often. But the attempt to use risk scores to target inspections of hospitals failed. Researchers found no link between the risk scores and the outcomes of subsequent inspections.

This is not to say that regulators should forego risk assessments. But they should be aware of the growing discussion about the reliance on risk assessment in policing and other fields. Critics and commentators have pointed out problems with errors and bias infecting the algorithms that are at the basis of the predictions and programs. There is also much talk about fairness and due process and transparency. Among other things, those who are affected by risk assessments likely have the right to know about them and to challenge them. To this extent, risk assessments are much the same as no fly lists or police watch lists. These issues and more will be following close behind as regulators move toward more sophisticated methods of assessing and controlling risk.

The information and comments in this article are for the general information of the reader and are not intended as advice or opinion to be relied upon in relation to any particular circumstances. For particular application of the law to specific situations, the reader should seek professional advice.

Interested in more articles like this one? Email publications@weirfoulds.com to join our Professional Self Regulatory Bodies Newsletter subscriber list.

WeirFoulds^{LLP}

www.weirfoulds.com

Toronto Office

4100 – 66 Wellington Street West
PO Box 35, TD Bank Tower
Toronto, ON M5K 1B7

Tel: 416.365.1110

Fax: 416.365.1876

Oakville Office

1320 Cornwall Rd., Suite 201
Oakville, ON L6J 7W5

Tel: 416.365.1110

Fax: 905.829.2035