Concepts and methods for strengthening risk cultures

Panel session on the Culture of Risk, LabWorks 2015, London
9 July 2015

Good afternoon. I am going to make four points that are intended to both raise a note of caution and point the way to a promising approach to fostering experimentation and innovation in the public sector.

The first point is that Governments are the uncertainty and risk managers of 'last resort' – they must handle a range of critically important problems and threats (including existential threats) that markets, and the companies whose capabilities shape those markets, cannot cope with. This means that experimentation and innovation in the public sector operates in an especially demanding environment – an environment in which the costs of negative unintended consequences can be truly dreadful (and put many lives and risk). The 'appetite for risk' and the failure rates encountered when innovating in the private sector do not face this unique challenge – which means that it is unwise to simply translate private sector innovation stances into a public sector context.

Indeed, when private sector innovation goes wrong - e.g. the regulatory system allows damaging innovations, such as a harmful drug, to be sold then it is governments that must intervene and try to pick up the pieces - and try to learn how to avoid such problems in the future. Of

course, experimentation and innovation are also key to helping governments to improve their role as uncertainty and risk managers of 'last resort'.

This is why governments try to build consensus-based 'immune systems' that attempt to identify possible negative unintended consequences - attempting to identify and pre-empt possible negative unintended consequences both to the nation's economy, society and environment and international relationships and obligations. Consequently, the risk-reward relationship for the public sector varies significantly from that of the private sector. Therefore, it is imperative to recognise these differences in order to enhance the effectiveness of efforts to stimulate public sector innovation.

In the private sector, the litmus test for attempts at innovation is market success. Not all innovations prevail in the market, and indeed various other factors mean that the 'best' solutions may not become the dominant solutions. However, markets do enforce selection processes that tend to eliminate less competitive solutions. Competing firms therefore do their best to second-guess what will prevail in the market, often applying vigorous structured decision-making processes (such as Stage-Gate methods) to weed out less promising concepts and solutions.

However, in the final analysis it is the market, and the social and cultural preferences that are reflected in markets, that will decide which innovations succeed and which do not.

In contrast, the public sector must draw heavily upon external and internal expertise to weigh-up complex risks of generating negative externalities. When there is no market-based 'shortcut' available the sheer weight of evidence that may need to be assessed poses major challenges, and raises important questions about whether 'hierarchies of evidence' are required to deal in a rational way with the sheer quantity and complexity of information available. The consequences of incorrectly judging what will and won't work when seeking to innovate are disproportionately greater for this type of public sector innovation than for private sector innovation.

Unfortunately, many governments forget this aspect and make the mistake of approaching their innovation capability as a 'deficit' (we aren't as good as the private sector) rather than in a more realistic manner aligned with their distinctive, and usually demanding role. This misconception has been encouraged by the 'new public management' ethos characterised by the privatisation of certain public services and a strong emphasis on the use of targets and performance measures to drive performance and demonstrate transparency and accountability.

So, my first point is that we should be wary of stances that frame public sector innovation in a 'deficit' manner when compared to the private sector. We need to do a better job at articulating why the public sector must approach the relationship between innovation and risk differently than the private sector – and in the light of governments' role as uncertainty and risk managers of last resort.

My second point is that there are ways of articulating this required agenda. It is possible re-frame innovation in the public sector as the primary means of coping with uncertainty and risk. We can do this by drawing on tools developed in the information and communication technology sector - tools that are spin-offs from path-breaking national security work in WW2, notably code-breaking and interpreting radar signals. These signal processing and machine learning tools (as they have now become) focus on the incidence of issues such as 'false positives' (a positive test result when there is no condition present) and 'false negatives' (a negative test result when there is a condition present) when analysing information. In other words they focus on the capabilities to try to *avoid* when drawing conclusions when information is (as it usually is) sparse and ambiguous. Signal processing and machine learning techniques usefully frame things as a 'confusion matrix': a 2 x 2 taxonomy that categorises test results as either true positives, false

positives, true negatives or false negatives. False positives and false negatives can paralyse or distort learning and adaptation.

Judgment calls made in such confusing circumstances are necessary - but risky. There is scope for 'worse than random' decisions to be made in the sense that conclusions are drawn that are negatively correlated with reality (a sort of Alice through the looking glass world of reversed interpretation). In such a situation, governments would have been better off making decisions randomly in the face of such confusing signals (this is a familiar and sometimes high-profile issue when judgment calls are made in security intelligence assessments). The ability to identify and diagnose how effectively the policy learning cycle functions in these terms may turn out to be particularly useful. So, there are the tools available. Indeed, using tools such as the 'confusion matrix' with its emphasis on false positives and false negatives can help the public sector to better execute the demanding role as uncertainty and risk manager of last resort.

My third point is that we should not hold back from making the conceptual leap to framing *every* government intervention as the tests of adaptive and evolving hypotheses with a binary test results (true or false). This is because it allows the potential for false positive and false negative test results to be articulated, which in turn allows us to put in place learning mechanisms for improving the accuracy of these tests.

These signal processing concepts are now playing a major role in shaping what doing evidence-based medicine involves (where avoiding false positive and false negative diagnoses is critically important). Indeed, the Randomised Control Trials (RCTs) approach that has been attracting such attention in the public sector can be thought of as one way of minimising the likelihood of arriving at false positive and false negative test results - but not always with as much success as the public policy community tend to assume. False positive findings are surprisingly common in clinical RCT studies.

The ability of a public sector organisation to understand its own diagnostic capability by examining it through the lens of these signal processing concepts, and to use this framework to set objectives for experimentation and innovation, can be really important.

My fourth and final point is very brief. It is that successfully articulating this 'binary' approach to the analysis of competing hypotheses using signal processing and machine learning tools opens up a rich new avenue for people working on the ground in delivering public services to suggest imaginative new hypotheses based on their accumulating experience and rich tacit knowledge – both diagnostic and as experimental interventions. The conducive and vibrant comparative testing environment that I am alluding to encourages this 'bottom' up approach by providing the concepts, methods and values needed.