You spend a lot of time teaching regulators. What do you teach them?

A number of things separate regulatory functions from all of governments' other functions. Social regulation largely focuses on the control of harms of one kind or another: crime, pollution, smuggling, occupational hazards, transportation hazards, diseases, food-poisoning, and so on. So regulators are more involved in containing or mitigating harms as the core objective, as opposed to delivering services. They also have at their disposal the coercive power of the state. They can deny licenses, put people out of business, seize assets, deprive them of their liberty, and in certain circumstances kill them – all in the name of public purposes. These are awesome responsibilities, and the menu of managerial techniques imported into government from the private sector is really not much help for regulators. Over the last twenty years, the tools offered and managerial methods available have mostly had to do with quality management, process improvement, process redesign, and excellence in customer service.

Regulation represents an anomaly in the context of customer service. The people regulators deal with are often not the people they want to make happy, and often not the ones paying for services. Regulators don't want to "delight their customers so they come back for more," which is the private sector motto. Regulators are expected to control risks, reduce harms – prevent bad things from happening. They need some relevant guidance on those aspects of their jobs.

Over the last five years or so, we've seen non-regulators starting to show up at my courses asking for advice and guidance. Some come from education, a lot from public

Unravelling Risk

Malcolm Sparrow is one of the world's top thinkers on government regulation. He authored the best seller *The Regulatory Craft* and as we go to press is about to publish a new work. He spoke with Vic Pakalnis of our editorial board.

interviewed by Vic Pakalnis

health, and increasing numbers from international non-governmental organizations and the not-for-profit sector.

When they come, I say to them, "I think you've come to the wrong course – why are you here? You're not a regulator." The answer they give is, "what you're teaching here is actually operational risk-control. And we have to do that, too, even though we are not regulators."

So I've ended up dealing with the United Nations and other NGOs on the control of various things. Some

groups want to deal with corruption, others terrorism, trafficking in women and children, or poverty alleviation. The not-for-profit organizations want to know what kind of contribution they can make, and foundations concerned about these issues want to know what kind of work to fund, and what kind of results to expect from the people they do fund. Although these bodies are not public agencies, they have similar kinds of questions: "How do I divide the task? What choices are involved? What meth-

ods should I consider, or not consider, for this? How will I know if what we did actually made a difference? How do we get the maximum impact for the money we spend?"

Would you have advice on how regulators who have adopted risk control methodologies might optimize them?

Let me clarify a few things about use of the word "risk." I'm worried about using the word in the context of risk-control because it means very particular things to some groups. For example, it has been used for many years in the financial sector, where the traditional notion of risk-management is about balancing risk and return in an investment portfolio. Others have used "risk" to mean risks that result from technology advances and human action, distinct from "hazards," which are naturally occurring.

Certainly there is now a burgeoning literature on risk management, risk analysis, risk communication and risk control. That literature tends to focus on small-probability events (at least, events which are small probability for any one individual) and pays a lot of attention to human psychology, perception and assessment of risks, and the behavioral consequences. We've learnt a lot from experimental psychology and behavioral economics about how people react to and treat different classes of risks.

In dealing with regulators, though, and focusing on the operational challenges of control, I prefer to undo the distinction between unlikely harms and well-established patterns of harm, such as those involving crime, pollution, and occupational hazards. We're not dealing just with things that might happen, we're dealing with things that happen often and that present established patterns. But potential or actual, all of these are bad things; I call them "harms" for want of a better word. In terms of the operational and analytic methods involved, I'm not sure there's a significant operational distinction between them.

So when we discuss the practical side of control operations, I like to undo the distinction – which is why I use the word "harms" a lot in the new book I am writing. The art of controlling them involves certain patterns of thought and action: the systematic disaggregating of big, broad generalities into specific, well-defined problems, followed by open-minded searches for tailor-made interventions that take significant bites out of the identified problem areas or risk-concentrations. **Discerning the parts or components of a problem, studying their structures and dynamics quite carefully, until you can see what it will take to unravel or "sabotage" them, and then skillfully undoing the parts, one by one – that's the art I try to teach, because that's the new professional practice that we observe emerging.** We also have to figure out all of the very complicated ramifications for organizational behavior – the role of analysis, the exercise of discretion at different levels of an organization. And agencies seem to need a lot of help understanding how to produce and present a new performance story that satisfies the public and legislative pressures on them. They really want a story that is a compelling and persuasive account of how they reduced risks, mitigated harms or eliminated problems.

Still the word "risk" for many people is a forward-looking, low-probability, might-or-might-not happen event. If you talk about an individual's risk of being killed on the highway in the next year, that's a small number. Their perception of that number might affect their willingness to drive rather than fly. But from the vantage point of those responsible for control, things look and feel different. If your job is to control highway accidents, this is not a low-probability thing. Suppose a state highway patrol department sees 400 fatal road accidents this year. They know for sure there will be between 300 and 500 next year, unless something dramatic changes. They just don't know who yet, even though they could tell you a lot about the victims' likely characteristics. The perspective from that level is not so concerned about psychology and perception of the risk itself, except insofar as it affects driver conduct and can be used instrumentally. Road fatalities, for the highway patrol, are an actual, well-established problem – like a "knot" that needs to be studied and then unraveled.

The term "risk" has too much particularity, and baggage. And that's what makes it problematic.

Tell me more about how we untie the "knots" and deal with society's problems?

First, I'm deliberately trying to embrace a broader range of bad things. Risk management literature tends not to talk about certain classes of risks. They don't talk so much about the risk of crime victimization, but they talk a lot about environmental risks, toxic risks, exposure to pathogens, and development of diseases. These are, in some ways, natural processes; there's no brain behind the risk.

The term "risk" has too much particularity, and baggage. **And that's what makes it problematic.**

Many people who study those systems and who are used to putting precise scientific numbers on a particular risk, find themselves befuddled when they understand we are talking about terrorists or thieves, smugglers or hackers, who are sitting in a room somewhere planning what they will do next. The "probability" of an explosion in a plane, for instance, can change dramatically, just because of a decision they took. In these contexts, trying to assess a probability in the normal scientific fashion, based on extrapolations from historical experience, means a lot less and will often get things very wrong.

So one thing I've been focusing on lately, and didn't do in *The Regulatory Craft*, is thinking about classes of harms that have particular properties. The risk literature has done that, to some degree. But it tends to look at classes of risks that affect human psychology and perception: "What makes a person willing to take a chance?" We talk about "familiar" versus "unfamiliar" risks,

and "voluntary" versus "involuntary" risks: why a person will complain about marginally polluted drinking water one day, but go sky-diving the next. Of course, we need to understand why humans behave the way they do, but I'm looking at these problems from the point of view of agencies or institutions (or sets of collaborators) who take on the task of mitigating societal risks for everybody. And I'm actually quite committed to the operational notion of reducing them, not merely transferring them or insuring them. Those are rather different policy issues.

One of the classes of risk I'm focusing on, for instance, is "invisible" risks: where the underlying level or scope of the problem is not known; where nobody knows how much to invest; where it is difficult to interpret performance metrics in an unambiguous fashion. Another group is those risks that have a brain behind them: where you're dealing with conscious opponents who display adaptive behaviours. This plunges you immediately into a game of intelligence and counter-intelligence, played against opponents determined to outwit you. This is an unfamiliar idea if you deal routinely with environmental or occupational hazards, which generally don't have a brain behind them searching for ways to kill or injure you.

I'm interested also in "catastrophic" risks, a very important subject these days. These are very small probability events that would have very serious consequences if they ever happened; where the actual number of incidents year after year is zero, zero, zero. Nevertheless you need to justify funding, decide how to allocate the budget, organize and delegate preventive
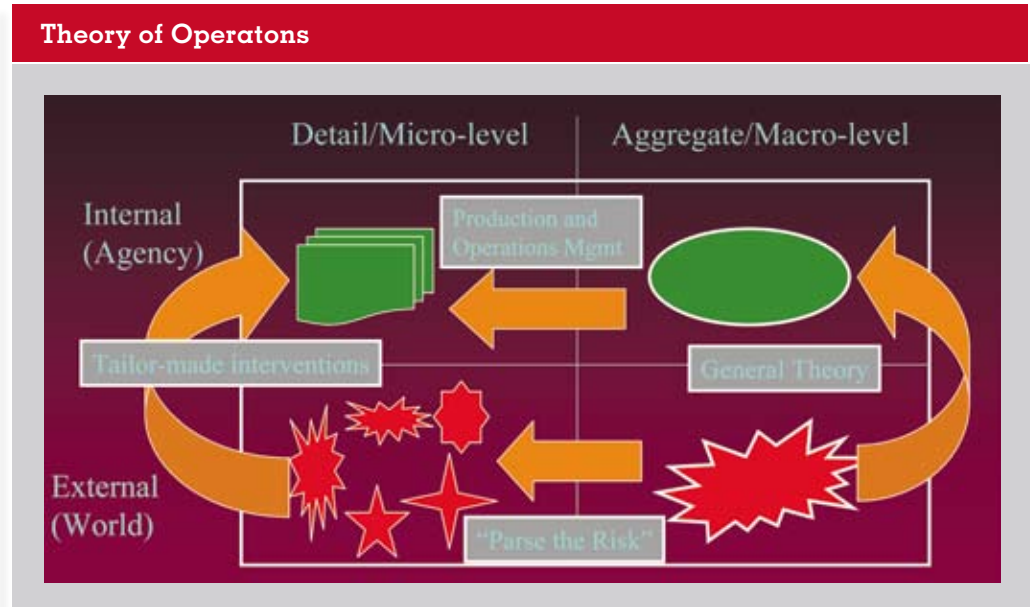
work, and show progress in driving the probability down further and further – even in the absence of the learning and practicing opportunities that more frequent but less-serious incidents would provide. People have a great deal of difficulty knowing how much to spend to keep something at zero, year after year, particularly when as far as the public is concerned it might have stayed at zero by itself. So: how to break up the work, how to allocate responsibility, how to measure progress and how to justify the expenditure for catastrophic risk prevention – these are all important and difficult subjects.

How would you guide regulatory agencies in trying to figure out how to deal with these very strange classes of harms?

I want them to be familiar with the most frequently observed properties of harms. As they confront different harms, they need to be quick to recognize when a particular harm is invisible, or involves conscious opponents, or is a harm in equilibrium that needs first to be dislodged before it can be guided into any other steady state, or is catastrophic, or is slow-acting (like cancer or ecosystem destruction), or is high-level (so it sits above the level of any one legislature or authority). Then at least they would be quicker to understand the consequences for the control challenge.

For instance, most environmental risks don't have a brain behind them. If you're protecting an eco-system or restoring lakes to fishable or swimmable conditions, there's no brain trying to find some other way to do environmental damage. If you're an environmental agency, you're not used to working with conscious opponents trying to deliberately outwit you. Then suddenly you encounter organized crime engaged in toxic waste-dumping, and you realize that they're really good at skirting new controls by studying agency operations and adapting quickly. And your scientists are flummoxed, because this "risk" now doesn't behave nicely. To deal with those kinds of opponents you need to think about intelligence operations, surveillance, use of undercover operations or informants or buying information – techniques basically unheard of in environmental protection, but well-known in other domains that deal with conscious opponents all the time, such as drug-smuggling, counterterrorism and military intelligence.

I'm eager to make these general notions broadly available, and clearly understood, so an environmental agency is quicker to recognize what they've just run into, and why and how it is different from the bulk of their work. It's that kind of awareness I'm trying to promote.

You use a two-by-two diagram to help public servants find solutions to big issues – could you elaborate?

It is a useful chart that illustrates the differences between two different kinds of organizational behavior. Following the "traditional" route, once you see a big, broad class of risk, you establish a big, broad class of solutions, with maybe four or five separate functional or programmatic approaches, and decentralized geographic structure overlaid. At one point in time, somebody figures out the "operational theory" – that

is, which methods they believe will affect the general class of harms, and that becomes the operational norm. They establish major programs, implement high volume core processes, and determine relevant functional expertise to acquire. Agencies seem to have a habit of doing that thinking once, and then operating on that basis for ten years or

more. These methods become ingrained models of behaviour for the organization.

What we are seeing more and more, and all around the world, is a very different behavior, regarded often as "innovative," in the regulatory fields. This alternate behavior involves keeping your eyes focused a little longer on the outside world, deliberately picking apart the generality to find very distinct concentrations (those are the smaller stars in the bottom left hand corner). Once these specific concentrations have been identified, they are then studied carefully, understood in their own right, and picked apart using tailormade methods that often were not in the ordinary or familiar toolkit for the agency.

This is where I use the analogy of finding knots, carefully studying them until you understand their structure and dynamics, then forming a

plan for undoing them. When agencies do act that way, they often get very substantial reductions with the specific knots they work on – not 10% or 15% reductions, but 70%, 90%, and sometimes complete elimination of a problem. The achievements under that model arise from major reductions in identified sub-components of a

risk, rather than across-the-board or distributed reductions at lower levels.

Presenting it in a two-by-two chart almost makes it look like a trivial operational distinction. It is, in fact, a completely different kind of behaviour inside large organizations; one that is extremely difficult to produce, maintain, manage and defend.

It involves extraordinary ranges of choices that are not normally made, about how big or small a problem to take on; in what dimensions the problems are best defined; how many to take on; whether to focus on scary long-term ones or smaller and shorterterm wins; whether to focus on projects that you can do by yourself, or whether to include some that require complicated partnerships. There is a portfolio management task here that requires a lot of judgment and subjective choices. Many agencies are not ready to make

such choices, and wouldn't be confident in their ability to defend them, either.

From watching some agencies get this right, and others get it quite wrong, some rather simple and powerful principles emerge that seem to be useful rules of thumb. The most important one is: respect the natural shape and size of the problem itself. Organize around that. Resist the temptation to squeeze the problem into your structure, to chop it up according to your organizational chart. Doing that is unnatural in terms of the risk itself, even if it's easier and quicker for you. It's a very ineffective approach of undoing the "knots." It's more effective if you can handle the organizational awkwardness and organize around the harm itself: do the thinking once, at the right level.

As agencies acquire the fluidity and flexibility to organize around harms, they eventually end up producing structures in-house that actually reflect the structures of the harms themselves. If you are dealing with one major risk, which has four sub-components, and one of those sub-components has two sub-sub-components, then chances are you should end up with four projects, and one of them will have two subprojects. That's the structure of the thing in the external world. You set up a structure so that each distinct object is tackled by one distinct group, which can tailor its strategy. And where a risk is composed of several unlike objects, you decompose it before tackling it.

Really, what you want for each particular problem is to have the thinking done once, by the right people, at the right level, and engaging all of the relevant knowledge and expertise.

Complex problems will require several different perspectives and collaborative actions; they normally don't fit neatly inside existing divisions of an organization. Nor do they usually align neatly with organizational boundaries. So organizing around the risks themselves is very awkward; it is just too hard for many agencies. But when you do, quite complex problems can be very substantially reduced. Public servants who have done this type of work tell me that it is the most challenging work they have ever done, but the most rewarding, too. They know that through these kinds of projects they actually made an important difference: saved lives, eliminated a hazard, cut in half an injury rate.

In *The Regulatory Craft*, you indicated that regulatory agencies should pick important problems – does your advice in your new book build on the same principle?

Yes. It contrasts the problemoriented methodology with traditional patterns of thought and behavior. But it hopefully serves a broader audience, beyond regulatory agencies alone. It also considers a broader range of classes of harms, such as corruption control, disease control and poverty reduction. Some of these are not normally regarded as regulatory functions. Nevertheless, they are chronic problems, and susceptible to the same patterns of intervention. I believe the audience is broader, and they don't need the backdrop of regulatory policy or history. They do need to know what's distinctive about controlling harms, and what kinds of obstacles they will likely confront in that type of work.

Some of these areas – like poverty alleviation – have the same kind of analytic traditions that afflict regulatory problems. Often the analysis has been traditionally dominated by economists, and involves regression or correlation analyses performed at macro-level, often at the level of nations. That kind of analysis shows us that "Oh! The level of corruption is or is not inversely correlated with public-sector salaries, or with the volume of international trade agreements, or something else." These macro-level observations are interesting, and sometimes useful, but quite often not actionable in an operational sense. They show us the various giant levers that might or might not be pulled. But many of those levers (like economic growth or trust in the justice system) are pretty hard to move and seem to have a mind of their own.

I'm interested in offering practitioners the opportunity to climb down from that level into the texture of the problem and understand that the corruption problem is not one corruption problem; it's probably seven different corruption problems, and each of those seven has different concentrations in different places, and is best expressed in quite different dimensions.

It's fun to watch each of these major risk areas discovering these underlying truths, one by one, at different times, and scarcely ever informed by all the other analogues available. Professor Jeffrey Sachs' latest book, *The End of Poverty: Economic Possibilities for Our Time*, says much of this concerning poverty control. He refers to a "clinical" approach to poverty reduction, which emphasizes individual diagnosis of each nation, understanding the shape, distribution and nature of its poverty problems.

This is very much in contrast with the tradition of the World Bank and other poverty researchers, who have defined poverty in terms of threshold household incomes (e.g., under $2.00/day adjusted for currency values); they apply that definition all across the world and study how we can reduce these numbers.

It's quite different, though, to say, let's think about poverty traps and what types of poverty traps there are: those induced by unemployment; by catastrophic health events; by bereavement; by fertility (too many children). There are poverty traps having to do with lack of empowerment in voting rights. And some poverty researchers (normally the people that do case-work and thick-description, and work a lot in the field as opposed to purely crunching numbers) are aware of these richer problems and individualistic textures lower down in the chain. People at the top, the macro-economists, deal only with the aggregates. People at the bottom do reactive work, incident by incident, operating response processes of one kind or another. All of this important work is neither at the highest level, nor at the lowest. It's all awkward, in-between – above the level of cases, below the level of generalities. In my view, that's where all the important work of controlling risks takes place. But it often seems elusive, amorphous, entirely voluntary and way too difficult to organize. I'm very much interested in structures and methods that one can use to work at these in between levels.

People need methods and guidance – to know the difference between a knot and a tangle. When is it one knot, and when is it three that you can deal with sequentially. Knowing how many to tackle, when and why; knowing when you have bitten off too much, or when you have bitten off something too small to be significant; how to construct your "portfolio" of bites into a particular harm so that you have some hope of making an overall difference.

Have you decided on the name of your new book?

We played with all kinds of names, including Undoing Harms, to capture the analogy with untying knots. But there's a difficulty with the word "undoing." It means not only "untying" but also "reversing." That chronological meaning, turning back the clock, suggests to some that you're only dealing with reaction after the fact, and that you are excluding preventive plans. That's absolutely not true, and so that language is problematic. I rather like The Character of Harms, with a subtitle that shows this is for practitioners facing the operational challenges of controlling them. So maybe The Character of Harms: Operational Challenges in Control. We'll see.

Dr. Malcolm Sparrow is author of The Regulatory Craft *and professor of public management, John F. Kennedy School of Government, Harvard University.* ⊞