Risk and Regulatory Policy
Improving the Governance of Risk

Chapter 6

Risk-based Regulation: Choices, Practices and Lessons Being Learnt

by
Prof. Julia Black, London School of Economics and Political Science, United Kingdom*

This chapter identifies key aspects of the risk-based frameworks of eleven regulators in four countries across four sectors: environment, food safety, financial markets and health and safety. Risk-based frameworks contain real choices as to the types and levels of risk the regulator is prepared to tolerate. Risk-based regulation therefore requires regulators to take risks. In practice the risk-based frameworks themselves have risks and a regulator's risk tolerance is ultimately driven by the political context. The chapter explores how these are addressed. Section 6.1 of this chapter defines risk-based regulation, explores the motivations for its adoption, sets out the main elements of risk-based frameworks, and provides some examples. Section 6.2 explores key questions that arise in practice with respect to each of these elements. Section 6.3 examines some of the main issues and challenges which have arisen in implementation. Finally, Section 6.4 discusses the evaluation of risk-based frameworks and identifies lessons learned.

* This chapter was prepared by Professor Julia Black, Director of Research, Department of Law and Research Associate, ESRC Centre for Analysis of Risk and Regulation (CARR), London School of Economics and Political Science.

Introduction

Risk-based frameworks are increasingly becoming seen as a necessary attribute of "better regulation". Risk-based frameworks enable regulators to channel their resources to those issues which pose the greatest risk to the achievement of their objectives. In their narrowest form, risk-based frameworks are used to allocate inspection resources. However for an increasing number of regulators, risk-based frameworks are being developed to help them structure choices across a range of different types of intervention activities, including education and advice.

Risk-based frameworks appear technical and mundane, but they contain real choices about what matters to that regulatory agency and what does not. The fundamental question in any risk-based regulatory regime is what types and levels of risk is the regulator prepared to tolerate. Regulators do not often articulate what their risk appetite is in public, or even private. Setting that risk tolerance can be an extremely challenging task. Better regulation enthusiasts usually emphasise the positive aspect of risk-based frameworks – that they require regulators to focus on their priorities. But risk-based regulation is a zero-sum game. Resources which are spent in one area are not spent somewhere else. The flip side of focusing on priorities is that regulators have to identify which risks or levels of risk they are *not* prepared to devote the bulk of their resources to preventing.

In making that determination, regulators are bound to make an error. Risk-based regulation therefore requires regulators to take risks. Regulators, and their political supervisors, have choice. Should they err on the side of assuming a firm does pose a risk when it does not (in statistical terms, a Type II error), or err on the side of assuming that a firm does not pose a risk when in fact it does (a Type I error). These choices have always been made implicitly within regulatory bodies. In risk-based systems, they are rendered explicit. The consequences are significant. If regulators err on the side of assuming firms are risky when they are safe, they run the risk of being accused of over-regulation, and of stifling business and innovation. If they err on the side of assuming firms' activities are safe when they are risky, they run the risk of failure. That failure, as the financial crisis demonstrates, can be far reaching.

In practice, a regulator's risk tolerance is ultimately driven by the political context. All regulators face political risk, the risk that what they consider to be an acceptable level of risk will be higher than that tolerated by politicians, the media and the public. For regulators, minimising political risk is often the overriding concern. The higher the political salience of a sector or risk, the less will be the regulators' tolerance of failure in that particular area. The political context is often fickle, however; issues that were not salient suddenly become so, and *vice versa*. This has consequences for the allocation of resources, which may not always go where the risk model says they should.

Risk-based frameworks also have other risks, notably model risk, that the model does not capture all the relevant risks, and implementation risk, that it is inadequately implemented. This chapter explores how these are addressed.

Aims and scope of the research

The purpose this chapter is to consider the development and role of risk-based approaches to regulation and to identify policy recommendations that can have broader application for risk-based regulatory strategies. The chapter identifies key aspects of the risk-based frameworks of eleven regulators in four countries across four sectors: environment, food safety, financial services and health and safety (see Annex 6.A1 for details). This chapter does not attempt to provide a systematic overview of the state of risk-based regulation in each of these areas, nor does it attempt to set out detailed comparisons. Instead, the chapter focuses on some of the key policy issues in the design and implementation of risk-based frameworks that have arisen in these areas. It looks at how regulators have addressed these issues by comparing the choices they have made in the design and operation of their frameworks. It also explores the different challenges involved in "doing" risk-based regulation and the experiences that regulators have had with its implementation. Throughout, this chapter is concerned with drawing out some of the lessons that can be learnt through examining these frameworks, with a view to informing policy recommendations for the development of risk-based frameworks by other regulators.

Outline of the chapter

The chapter is divided into four main sections. Section 6.1 defines risk-based regulation, explores the motivations for its adoption, sets out the main elements of risk-based frameworks, and provides some examples. Section 6.2 explores the issues that regulators have found arise in the design of risk-based frameworks. Section 6.3 examines some of the main issues which have arisen in implementation. Section 6.4 discusses the evaluation of risk-based frameworks and identifies key challenges and lessons learnt.

6.1. What is risk-based regulation?

Risk-based regulation is a relative newcomer to the lexicon of regulation. It can be used to refer to anything from a loose agglomeration of approaches expressed in terms of risk, to highly structured and systematised decision making frameworks (see also Hutter, 2005). It is usually given one of three broad meanings. The first refers to the regulation of risks to society: risks to health, safety, the environment, or less usually, financial well-being. In this respect, "risk-based" regulation has long been used by regulators and legislators to determine whether or not an activity should be regulated, or what level of preventive measures firms or others should take.

The second meaning, which is particular to banking and insurance regulation, is a far more specific one: it is the use of firms' own internal risk models to determine the amount of capital banks should set aside. This model of "risk-based regulation" is entrenched in the Basle II capital adequacy rules, and enacted in the EU by the Capital Requirements Directive.

The third meaning of risk-based regulation is that on which this chapter focuses. It refers to the use of systematised frameworks of inspection or supervision which are primarily designed to manage regulatory or institutional risk: risks to the agency itself that it will not achieve its objectives. In this third sense, risk-based regulation involves the development of decision-making frameworks and procedures to prioritise regulatory activities and deploy resources, principally relating to inspection and enforcement, based on an assessment of the risks that regulated firms pose to the regulator's objectives.

In risk-based approaches, the focus is not on the potential risks that individuals or the market economy may face from the actions of firms *per se*, but on the risks the regulator faces in failing to achieve its objectives. The objectives of the regulator are translated into a rubric of risk, and their focus becomes the attainment of those objectives. Risk-based regulation thus requires regulators to explicitly define their regulatory objectives, and to translate their statutory mandates into operational objectives. Whether or not the regulator translates the objectives in a way which is supported by the wider polity remains an open question, however. It may be that there is congruence between the two, but this cannot always be assumed. In practice, there is often a misalignment, and regulators are driven by changes in the political and social context to address risks that they might otherwise have regarded as low priority.

It is fair to ask to what extent is the current flurry to develop "risk-based" approaches simply the dressing up of old systems and processes in new, more fashionable clothes? In Meyer and Rowan's familiar argument, organisations adopt structures and follow procedures not just, or not even, to achieve goals, but to gain legitimacy in the widest sense (Meyer and Rowan, 1977). The rhetoric of "risk management" and "risk-based" approaches combines a sense of strategy and control in a way which is politically compelling; moreover, framing one's actions as "risk-based" is, in the current climate, a useful legitimating device. But the framing of the regulatory task in terms of risk has the potential to have more than a rhetorical effect: it imports particular conceptions of the problem at hand, and leads to the framing of a solution in a particular way. Most notably, "risk-based regulation" introduces a matrix of assessments which focuses not, or not only, on economic costs and benefits, but on uncertainties, impacts and probabilities (Black, 2005a).

Risk-based frameworks contain real choices about what matters to that regulatory agency and what does not. For they require regulators to identify what risks or levels of risk it is *not* prepared to devote the bulk of its resources to preventing. We are familiar in debates on societal risk regulation of the choice between Type I and Type II errors: of erring on the side of caution (assuming something is risky when it is not) or erring on the side of risk (assuming that something is safe when it is not) (Schrader-Frechette, 1991). The debate usually operates at the level of deciding whether an activity should be regulated or not: in writing rules or setting standards, should regulators err on the side of protecting consumers (making Type I errors) or favouring producers (making Type II errors). However, regulators also face the same choices in their own organisational decisions of what level of attention to give to any one firm. The consequences are significant. If they err on the side of assuming firms are risky when they are safe, they run the risk of being accused of over-regulation, and of stifling business and innovation. If they err on the side of assuming firms are safe when they are risky, they run the risk of failure. The experience of the UK FSA in it supervision of Northern Rock, and indeed the credit crisis more broadly, is an excellent example (FSA, 2008). The FSA assumed the bank's business model was safer than it was; but intervention any earlier would have created political resistance on the grounds that they were interfering in a highly profitable business.

The key motivations for adopting risk-based approaches

There has been a significant increase in the use of risk-based frameworks for inspection and supervision in a range of countries and across a number of sectors, by both state and non-state regulators (see Black, 2005a; 2005b; Hutter, 2005; IOPS, 2007; Brunner *et al*., 2008; Rothstein *et al*., 2006).

Although the precise reasons for each regulator to adopt a risk-based approach are obviously unique, both the research done for this study and the findings of other studies suggests that there is a common core of motivations (Black, 2005b; IOPS, 2007; Hutter and Lloyd Bostock, 2008; Rothstein *et al.*, 2006). These are broadly functional, organisational, environmental (in the broadest sense), political and legal.

First, regulators have turned to risk-based frameworks in an attempt to improve the way in which they perform their functions. They have adopted risk-based frameworks in an attempt to facilitate the effective deployment of scarce resources and to improve compliance within those firms which posed the highest risk to consumers or the regulators' own objectives. Risk-based frameworks are also adopted to improve consistency in supervisors' assessments of firms, to enable regulators with broad remits to compare risks across a widely varying regulated population within a common framework. More broadly, risk-based frameworks are being adopted part of a more general desire by regulators to become more "risk aware" and less rule-driven in their activities.

Second, risk-based frameworks have been adopted to address a range of internal organisational concerns. In particular, they have been introduced to provide a common framework for assessing risks across a wide regulatory remit, and to deal with mergers of regulatory bodies. They have also been seen as a way in which to improve internal management controls over supervisors or inspectors. In federated structures, where the regulatory regime is split between central government and local authorities or municipalities, risk-based frameworks are also used to provide a framework for central government control. An example here are the risk-based frameworks for inspection issued by the UK Food Standards Agency with which local authorities in England and Wales have to comply.

Third, risk-based frameworks have been adopted in response to changes in the market and business environment. For example, banking regulators started developing risk-based systems in tandem with an increasing preoccupation within banks in using risk-based assessments for their own internal purposes. Food regulators in the US point to the introduction of HACCP as facilitating the introduction of a risk-based inspection system (FSIS, 2007).

Fourth, the political context can be highly significant. Risk-based frameworks have been adopted in response to previous regulatory failures, and to provide a political defence to charges of either over- or under-regulation by politicians, consumers, the media or others (Black, 2005a; 2005b). More generally, having a risk-based framework has increasingly become a badge of legitimacy for a regulator. Risk-based systems are a key part of the "better regulation" framework, and as such are a core attribute that regulators need to possess.

Finally, as risk-based regulation becomes seen as a functionally efficient tool for improving better regulation, politicians and others are increasingly requiring regulators to adopt such frameworks by law. In the area of food safety, for example, EC regulations require that inspections be carried out on a "risk basis" (EC 882/2004). In the UK, regulators are now subject to new statutory duties of "better regulation" set out in the Compliance Code. These include the requirement to adopt a risk-based approach to inspection (DBERR, 2007).

The main elements of risk-based approaches to regulation

The frameworks vary considerably in their complexity. However all have a common starting point, and four common core elements.

The key element of risk-based frameworks for allocating resources is that the starting point is risks not rules. Risk-based frameworks require regulators to begin by identifying the risks that it is seeking to manage, not the rules it has to enforce. Regulators are usually over-burdened by rules. They cannot enforce every one of these rules in every firm at every point in time. Selections have to be made. These selections have always been made, but risk-based frameworks both render the fact of selection explicit, and provide a framework of analysis in which they can be made.

The frameworks themselves have four core elements. First, they require a determination by the organisation of its own risk appetite – what type of risks is it prepared to tolerate and at what level. This can be an extremely challenging task for a regulator. In practice, a regulator's risk tolerance is often ultimately driven by political considerations. All regulators face political risk, the risk that what they consider to be an acceptable level of risk will be higher than that tolerated by politicians, the media and the public. Political risk is in practice a critical element in any risk-based system, as discussed below.

Second, risk-based frameworks involve an assessment of the hazard or adverse event, and the likelihood of it occurring. Terminology varies: food and environmental regulators tend to talk in terms of hazards and risks; financial regulators talk in terms of impact and probability. Two broad categories of risk are identified: the inherent risks arising from the nature of the business's activities, and in environmental regulation, its location; and management and control risks, including compliance record. These assessments may be highly quantitative, or be mainly qualitative. The methods by which management and control risks are combined with or offset against inherent risk scores varies, but broadly speaking management and controls can either exacerbate the inherent risk or mitigate it.

Third, regulators assign scores and/or ranks to firms or activities on the basis of these assessments. These may be broadly framed into three categories or traffic lights, or there may be a more granular scoring system, with five or more categories.

Fourth, risk-based frameworks provide a means of linking the organisation and of supervisory, inspection and often enforcement resources to the risk scores assigned to individual firms or system-wide issues. In practice, resources do not always follow the risks in the way that the framework would suggest, however, as discussed further below.

The tables below briefly summarise some of the risk-based frameworks in the different sectors covered by this research. All the risk-based systems investigated are outlined in Annex 6.A3.

Table 6.1. **Financial services: comparison of the risk-based frameworks of FSA, APRA and DNB**

Organisation Element	FSA: Arrow (Advanced Regulatory Risk Operating Framework)	APRA: PAIRS (Probability and Impact Rating System)	DNB: FIRM (Financial Institutions Risk Analysis Method)
Date first introduced	2001 (Arrow 1); 2006 (Arrow 2).	2002.	2006-07.
Outline of risk assessment framework	Risk = impact of risk × probability of risk occurring.	Assess inherent risk and management and control to derive net risk. Then consider capital support to determine overall risk of failure.	Inherent risk minus management and control = net risk. Net risk minus capital support = overall risk of failure.
Risk scoring and categorisation	Four × four matrix. Impact L, ML, MH, H. Probability L, ML, MH, H. On site risk assessment and relationship management for firms of ML impact and above.	Individual risk assessments prepared for all licensed entities. Two stage categorisation: ● Impact analysis based on asset size. ● Probability analysis based on scoring between 0-4 of key risk categories.	All institutions have individual risk analysis. Traffic light system (red for the highest risk; orange for medium risks; green for low risks).

Table 6.1. **Financial services: comparison of the risk-based frameworks of FSA, APRA and DNB** *(cont.)*

Organisation Element	FSA: Arrow (Advanced Regulatory Risk Operating Framework)	APRA: PAIRS (Probability and Impact Rating System)	DNB: FIRM (Financial Institutions Risk Analysis Method)
Risk identification – categories of firm specific risks	53 risk elements which are consolidated into 10 risk groups: ● Environmental. ● Customers, products and markets. ● Business processes. ● Prudential (credit, market, operational, insurance and liquidity). ● Customers, products and markets controls. ● Financial and operating controls. ● Prudential controls. ● Control functions (internal audit, enterprise-wide risk management and compliance). ● Management, governance and culture. ● Capital and liquidity.	● Board. ● Management. ● Risk governance. ● Strategy and planning.[1] ● Credit risk.[1] ● Market and Investment risk.[1] ● Insurance risk.[1] ● Operational risk.[1] ● Liquidity risk.[1] ● Capital support.	Inherent risks: ● Financial risks. ● Liquidity risks. ● Insurance risks. ● Operational risks. ● Integrity risks. ● Strategic risks. ● Governance risks.
Risk assessment against regulatory objectives	Each risk assessed against each of seven "risks to objectives" (RTOs) derived from its 4 statutory objectives (customer protection, market confidence, reducing financial crime and promoting public understanding): ● Financial failure. ● Misconduct/mismanagement. ● Consumer understanding. ● Fraud/dishonesty. ● Market abuse. ● Money laundering. ● Market quality.	Assessed against single objective: ● Financial failure.	Assessed against single objective: ● Financial failure.
Regulatory response	Risk mitigation programme.	SOARS (Supervisory Oversight and Response System) linked to PAIRS risk assessment.	Specified supervisory menus linked to risk score.

1. Assess both inherent risk and management and control for each category.

Table 6.2. **Environment: comparison of the risk-based frameworks of EA, EPA and IGAOT**

Organisation Element	Environment Agency (England and Wales)	Irish Environmental Protection Agency	Portuguese IGAOT
Date first introduced	2002, latest version 2008.	2007-08.	2009-
Outline of risk assessment framework	Probability and hazard analysis with respect to each attribute.	Probability and hazard analysis with respect to each attribute.	Probability and hazard analysis with respect to each attribute.
Risk scoring and categorisation	Individual detailed Opra analysis for bespoke permits only scores.	Individual assessment for all licensed activities/installations. 3 grade scoring system A (high) – C (low); each grade subdivided (A1-3; B1-3; C1-2).	Individual assessment for all IPPC (integrated pollution and prevention control legislation) activities/installations. 3 grade scoring system: (high, medium, low).
Risk identification – risk attributes	5 risk groups: ● Complexity ● Emissions and inputs ● Location ● Operator Performance ● Compliance rating using compliance classification scheme.	5 risk groups: ● Complexity. ● Emissions and inputs. ● Location. ● Operator management. ● Enforcement record.	5 risk groups: ● Complexity. ● Emissions and inputs. ● Location. ● Attitude of operator to the environment and sustainability of the attitude ● Compliance behaviour.
Risk assessment against regulatory objectives	Used with respect to emissions and waste management; anticipated for water quality discharge consent regime in 2009-10.	Used with respect to emissions, waste management and discharges into water/sewers.	Planned introduction in 2009 to emissions, waste management and discharges into water/sewers.
Regulatory response	Supervisory discretion.	Supervisory discretion.	Supervisory discretion.

Table 6.3. **Food: comparisons of the risk-based frameworks of the UK Food Standard Agency and the Food Safety Authority of Ireland**

Organisation Element	Food Standards Agency (England)	Food Safety Authority of Ireland
Date first introduced	1995, latest version 2008.	2000, latest version 2006.
Outline of Risk Assessment Framework	Hazard and impact analysis of activities.	Hazard and impact analysis of businesses.
Risk scoring and categorisation	5 categories A (high) – E (low).	3 categories (high-low).
Risk Identification/Risk attributes	Food hygiene: ● Potential hazard (type of food and method of handling; method of processing; number of consumers at risk). ● Level of current compliance. ● Confidence in management/control procedures. ● Specific risk assessment of potential contamination by specified micro-organisms.	Pre-populated score sheet scoring types of businesses. Businesses not listed to be assessed on basis of analogy with existing categories; and in addition: ● consumer profile; ● scale of the operation; ● type of food; ● nature of handling/processing; ● structure and layout of premises; and ● control systems.
Risk assessment against regulatory objectives	Food safety and public confidence.	Food safety and hygiene.
Regulatory response	Intervention scheme linked to risk levels; minimum levels of interventions (not limited to inspections).	Minimum levels of inspection set for each risk category.

Table 6.4. **Health and safety: the UK Health and Safety Executive's Field Operations Directorate framework for non-hazardous activities**

Organisation Element	Health and safety executive
Date first introduced	1990s; latest version 2008
Outline of Risk Assessment Framework	Analysis of risk, probability and nature of harm
Risk scoring and categorisation	Gap: Gap between level of risk firm is at and where it should be if in compliance. 4 categories of risk gap: extreme, substantial, moderate and negligible; 6 point rating scale for individual risk elements.
Risk identification/Risk attributes	Risk elements: ● consequences; ● likelihood; and ● extent. Categories: ● Safety. ● Health. ● Welfare. ● Competence and attitude of management.
Risk assessment against regulatory objectives	Health, safety and welfare.
Regulatory response	Supervisory discretion in line with enforcement management model.

6.2. Designing risk-based frameworks

The development of risk-based frameworks follows the pattern of many innovations (Black, Lodge and Thatcher, 2005). There have been a few "early adopters", and over recent years the number of regulators adopting some kind of risk-based approach has steadily increased. The later adopters have been directly or indirectly helped by the "early adopters". Regulators have communicated the detail of their frameworks and their experiences to other regulators through transnational networks, such as IMPEL in the environmental context, or by bilateral interchanges (see also Black, 2005b). Models of risk-based systems are thus spread across regulators, and modified each time. For example the Australian Prudential Regulation Authority's (APRA's) risk-based model has been adopted in modified form in a number of different countries. Regulators often "mix" models – so the Portuguese environment regulator, IGAOT, used a mixture of the Irish Environmental Protection Agency's framework, with that of the Dutch environmental regulator, VROM. The Irish EPA's framework itself drew on that of the Environment Agency

for England and Wales, and its food hygiene framework draws on that of the Food Standards Agency's Code of Practice for England.

Despite these patterns of learning, no two risk-based systems are identical in their form, and often differ significantly in their operation, even if in form they may have strong similarities. Some of these differences stem from their widely differing remits, and their location within governments. But others reflect strategic choices. As a result, the differences can be revealing. Risk-based frameworks are not neutral, technical instruments. Each aspect of a risk-based framework involves a complex set of choices. They require decisions by the regulator as to what level of risk or failure it is prepared to accept; what risks it will identify as requiring attention; what indicators and methods it will use to assess those risks, and how it will deal with the majority of firms that fall into the "low risk" categories. This section considers each of these choices in turn.

Risk tolerance

The fundamental question in any risk-based regulatory regime is how much risk is the regulator prepared to tolerate. Regulators do not often articulate what their risk appetite is in public, or even private. Those that have stated their risk tolerances publicly differ significantly between sectors. The financial regulators adopt, in theory, a non-zero failure policy, following the FSA's statement of this position (FSA, *Reasonable Expectations*). In the paper *Reasonable Expectations* the FSA noted there was a gap between public expectations of what regulators should or should not be able to achieve, and what "reasonable" expectations should be (FSA, 2003). The paper made it clear that "non-zero failure" meant that the regulator would not, and should not be expected to, prevent every "negative event": every financial failure of a firm, every incidence of non-compliance, every incidence of market failure, and that public and political expectations of what regulation can achieve should be modified in accordingly.

In food regulation, in contrast, the policy with respect to food additives and residues of pesticides and veterinary drugs is one of "notional zero-failure", although for contamination by micro-organisms, however, food regulators tend to adopt a standard of "as low as reasonably practicable". As a review of food regulatory systems observed, however, given the difficulties in obtaining reliable data and the public expectation that food should pose no risk, targets are usually defined in relative terms (a reduction of 25% over 2 years) rather than absolute terms (Slorach, 2008). Health and safety regulation in the UK also provides for a residual level of risk to remain, even when there has been full compliance with the rules. The requirement is that health and safety be assured "so far as is reasonably practicable".

Whatever their policy, and whatever their legislative framework, risk-based regulation requires regulators to take risks. This is extremely challenging for a regulatory organisation. They have to choose which risks or levels of risk are they *not* prepared to devote the bulk of their resources to preventing. As noted above, they have a further choice. In making that determination, should they err on the side of assuming a firm does pose a risk when it does not (in statistical terms, a Type II error), or err on the side of assuming that a firm does not pose a risk when in fact it does (a Type I error). These choices have always been made implicitly within regulatory bodies. In risk-based systems, they are rendered explicit.

In practice, the political context is determinative. The higher the political salience of a sector or risk, the less will be the regulators' tolerance of failure in that particular area. Indeed, several regulators deliberately calibrate their risk models in terms of their ability to maintain public confidence in themselves and in the sector they are regulating

(see Section 6.3 below). The political context is often fickle, however; issues that were not salient suddenly become so, and *vice versa*. This has consequences for the allocation of resources, which may not always go where the risk model says they should. Rather they go to the area which is most politically sensitive. As the current credit crisis illustrates, even a non-zero failure policy can be abandoned when the political, and systemic, stakes are too high.

Risk identification

Which risks?

The foundation of any risk-based approach is the risks on which it focuses. There is a multitude of risks on which regulators can focus, and regulators have to be selective. Clearly, in addressing model risk, the risk that the risk-based framework does not focus on the relevant risks, regulators have to choose these risks carefully. Regulators have taken different approaches to identifying and selecting these risks.

The starting point is usually the regulators' statutory objectives. The UK Financial Services Authority, for example, has framed the groups of risks on which it focuses as "risks to objectives". Lack of clear statutory objectives can thus be a hurdle to formulating RBFs. In the UK, the previous regime for pensions' regulation was hindered, amongst other things, by a lack of clear statutory objectives. Changing the legislation to introduce clear objectives thus facilitated the development of The Pension Regulator's (TPR's) risk-based approach.

A highly complex legislative framework can also be a hindrance. Regulators are often charged with implementing a significant number of individual pieces of legislation. The Netherlands environmental regulator, VROM, for example, is charged with over 270 specific legislative tasks. A key stage in developing its risk-based framework was therefore synthesising these into four different types of impacts of the activities which they were charged with regulating: health, sustainability, safety and social elements. Over time they have further refined their work into four work programmes: water, soil, safety and air quality, and now examine each type of impact with respect to each work programme.

Other legal duties can be also be relevant for identifying risks. The environmental regulators in the UK, Ireland and Portugal include those emissions to air and discharges to water and sewers which they are required to report to the European Environment Agency. They are required to collect this information, and so it makes sense to include it in their risk-based frameworks.

Public perceptions and expectations of the regulator can also be important. The Food Standards Agency in England and the UK Health and Safety Executive, for example, take into account the perception of public attitudes to risk in identifying which risks they should focus on, as well as their statutory objectives. In Ireland, the Irish EPA has included odour as a risk on which it should focus, as this gives rise to considerable complaints and can be resource intensive to deal with.

Risk-based frameworks can become highly complex as the number of risks on which they focus increases. The risk-based frameworks in use vary considerably in their degree of complexity. Those with simpler systems are often regulators who are just drawing up their risk-based systems, such as IGAOT, and/or whose regulated population is engaged in less complex or hazardous activities, such as the Irish EPA.

Finally, the amount of data that the regulator currently has can have a significant bearing on which risks they focus on. Regulators can only identify risks that they already know about. There is more over the danger that regulators only identify risks that they are confident they can manage, and leave other risks out of consideration. There was a

suggestion from some regulators that limitations on the data that the regulator currently had or thought that it could collect restricted the risks which it included in its risk-based framework. In some respects this makes sense: there is little point in the regulator trying to manage risks when it does not have the information on which to assess them. However, an incomplete set of risks does enhance model risk – the risk that the model itself is flawed. This will be returned to below in Section 6.3.

To address the problems caused by not having the right type of information, others who are just beginning to implement risk-based systems have deliberately designed the system in such a way that it maximises the amount of data they will receive. The UK's Office of Fair Trading (OFT), for example, wanted a risk-based system that would achieve three aims: it would enable the OFT to gain more information on the firms that had licences; it would achieve consumer protection objectives by weeding out those who already had consumer credit licences but did not really need them; and it would emphasise areas that the OFT already knew were high risk, notably debt management and debt guidance. The OFT has thus designed a system which requires licence applicants to give a significant amount of information, on the basis this will give it the data which it needs in order to refine its risk-based system in the future.

Risk indicators

Having identified the risks, regulators have to determine what the risk indicators should be. Risk indicators are those activities or events that are likely to result in the risk crystallising. Risk crystallisation occurs when the adverse event that the regulator was trying to prevent in fact happens; that there is a discharge, accident, food contamination, or financial failure, for example.

With respect to the processes for identifying risk indicators, the interviews revealed that some regulators use external consultants and those who were able to draw on others' frameworks had borrowed heavily from them. All the regulators interviewed, however, held intensive internal discussions as to what the relevant risk indicators were. These discussions drew on the knowledge of inspectors and supervisors as to the causes of previous failures or accidents. Tacit knowledge across the organisation as to what the warning signs were of impending failures was pulled together and then transformed into explicit risk indicators against which risk, in particular probability, would be assessed. In nearly all cases, the risk indicators derived from this process of internal discussion and distillation were judged by those developing and implementing the framework to have been more valuable than those proposed by external consultants.

Balance between objective and subjective indicators and assessments

The choice of risks and risk indicators is subjective, but the frameworks vary considerably in the extent to which the indicators they use can be assessed objectively or subjectively. The indicators used by the environmental and food regulators, for example, tend to be objective, quantitative measures. For the environmental regulators, complexity of the site is based on the types of activities carried out. The activities are defined objectively and grades defined based on judgements as to their significance, or legislative requirements. They are assessed with respect to indicators specifying the types of activities conducted, the capacity for production (not actual production), and/or the area over which the activity occurs. Thus, one of the indicators of complexity for the Irish EPA is the:

a) Production of non-ferrous crude metals from ore, concentrates or secondary raw materials by metallurgical, chemical or electrolytic processes.

b) Smelting, including the alloyage, of non-ferrous metals, including recovered products (refining, foundry casting, etc.) with a melting capacity exceeding 4 tons per day for lead and cadmium or 20 tons per day for all other metals: score is the mid-range grade G3 (POE, 3.4.1).

Emissions and discharges are assessed as capacity for emissions of particular substances measured by number of kilograms per year.

Quality of management is assessed on a yes/no basis. So in the Irish EPA framework, a low grading is given if the site has an approved environmental management system in place, a training plan and an environmental committee that meets regularly, combined with a low number of incidents reported. Enforcement history is quantified. A total score is given for the number of complaints received about the facility by the EPA, non-compliance notifications issued by it, the number of section notices issued and the number of convictions, all in the last year.

Other regulators include indicators which are assessed more subjectively. For the financial regulators, assessment of management, governance and culture, of control functions, and risks arising from dealing with customers are assessed on a qualitative, not a quantitative basis, as are risks to the firm from the external market environment. In the UK Food Standards Agency's framework, the assessments of hazard are based on the type of food, the nature of the handling, the type of processing methods used and the number of consumers at risk. Each is defined briefly, and the assessment criteria are deliberately framed in broad terms to encourage environmental health officers to develop and use their own professional judgement. Confidence in management and controls is defined in terms of the business's compliance record, and the likelihood of this being maintained at current levels, but again no quantitative inputs are used. The UK Health and Safety Executive also uses qualitative assessments of risks to health, safety and welfare and of management and controls in its risk-based framework for non-hazardous activities.

One mode of assessment is not necessarily better than another, and certain risks can be more easily assessed using quantifiable methods than another. However, the extent to which the risk-based framework uses qualitative assessments or relies almost entirely on quantitative assessments does have significant implications for the management, organisation and governance of the risk assessment process. This will be returned to in Section 6.3 below.

Risk assessment

Impact and probability

One of the critical issues in the design of a risk-based system is the relative role played by assessments of probability and impact or hazard. A bias towards impact means that regulatory attention is focused more on activities or events which have a relatively high impact but low probability; a bias towards probability means the regulator focuses more on high probability but relatively low impact events or activities. The regulators take quite different approaches to how they assess impact, the relative weights given to impact and probability, and the relationship between them. The choice is a political one, and the difference can be significant.

Impact measures. Impact is usually an assessment of the impact on the beneficiaries of regulation, broadly defined: the environment or consumers. So environmental regulators look at the maximum capacity of an installation to pollute, or discharge waste. Financial

regulators look at the size of the firm or fund as a proxy for impact measures. In the context of food regulation, proxies are the number of consumers and their nature (for example children or the elderly).

An approach used by several regulators with a high number of regulated firms is to use impact measures to divide the regulated population into risk groups in order to determine the depth and complexity of the risk assessment that will be applied to them. This approach is used by the Environment Agency in England and Wales, and financial and pension regulators in the UK, the Netherlands and Australia. Often, the initial categorisation is used to determine which groups will be subject to a risk assessment at all, and which will not. The group of firms that is subject to a risk assessment is then further subdivided into two groups – those who receive a simplified assessment, and those that receive a full assessment.

For example, under the new environmental permitting regime, the Environment Agency is moving to a system which divides licence holders into three groups or tiers. Tier 1 licences are for low impact activities, such as carrying household waste or fishing. Licence holders are simply required to pay for a licence, and there is a minimal level of random inspection carried out, principally for the purposes of protecting the integrity of the licence regime. Tier 2 licences are standardised permits and licences to which general binding rules apply. A simplified version of the risk framework, "Standardised Opra" applies to these sites. They are given a standard baseline score for four of the Opra risk attributes for each sector. That score is modified at a site level by the site's compliance score. The full Opra risk assessment applies only to Tier 3 licence holders. These are the more complex sites which are given bespoke permits, and the full, individualised version of Opra applies.

APRA uses a similar basis to categorise the pension funds which it supervises, with the smaller funds subject to a simplified version of its risk-based system. The UK Financial Services Authority has three versions of its risk-based framework. It has a "small firms" model which applies to low impact firms; these are not subject to individual risk assessments. Most of those in the medium-low impact categories are subject to ARROW light, which is a reduced scope risk assessment which focuses on core areas and sectorally important issues only. Medium high and high impact firms are subject to the full ARROW process, as are medium-low impact firms with a high probability (FSA, 2006).

Impact measures can also be designed or adjusted for more political ends, and to address political risk. This can be done explicitly, as in VROM's framework. Here "social impacts" are a separate category of impact, and are essentially there to capture issues which have current political and media salience. The UK Pensions Regulator also explicitly takes account of political risk in determining its impact measures for its higher risk firms. It defines impact of pension funds initially in terms of the number of members This gives two groups, those with over 1 000 members (1 600 schemes) and those with less than 1 000 members (83 000 schemes). The latter are a low priority, and regulatory action is focused mainly on education and guidance of trustees and members. The largest schemes are divided again into two groups: the 150-300 firms which pose the highest risk, and the next 300-1 600 schemes. High risk is defined in two ways: first, in terms of the number of members; second, in terms of the impact on the regulators' own reputation and future effectiveness. In other words, risks are identified the basis that if TPR did not pick these up it would be seen as a failure and the public would lose confidence in the pensions system

as a result (TPR, 2006). The regulators' political risk is thus clearly incorporated into the impact measure. Political impact can also be incorporated less explicitly. When Arrow 1 was first introduced, the FSA's impact measures for credit unions were deliberately inflated as the regulator wanted to ensure that they were given more regulatory attention than the scoring system would otherwise have permitted (Black, 2005a).

Others, notably the Food Standards Agency and the Health and Safety Executive, do not undertake this initial categorisation by impact. Moreover, impact measures focus on the scale of the harm, not its nature. There are some frameworks which include an assessment of the nature of the harm as well as its impact. The UK Food Standards Agency's framework, for example, includes both the nature of the micro-organisms that are present in the food and the number of people likely to consume the food. The Health and Safety Executive also combine consideration of the nature of the harm, the probability of it occurring and the number of people likely to be affected in their framework.

Focusing on the nature of harm can move impact measures away from an aggregate measure (how many, how much in total across an area/population) to a focus on individual impacts. The UK Office of Fair Trading focuses more on the nature of the impact on individuals than on the number of individuals that would be affected in its risk-based framework. Thus it identifies home debt collection as high risk, partly because it affects a significant number of those taking out consumer credit, but also because the nature of poor practices in home debt collection often involves violence and intimidation. Consumers are thus particularly vulnerable in these circumstances, even though the aggregate impact might not be great. Secured sub-prime lending is rated as high risk on similar grounds: that mis-selling of secured sub-prime lending will lead to default, which has a significant impact on the consumer. Indeed this example illustrates very well the difference between the financial and, to an extent, the environmental regulators' systemic approach to risk categorisation, and the individualised-consumer focus of the OFT's framework.

Relative weights and relationship of impact and probability. Impact measures are thus often used to determine when a full risk assessment should occur. Within that risk assessment process, probability and impact have different roles and are combined in different ways. They are also differently classified. The environmental regulators classify the inherent risks of a site as hazards, and compliance and management practices as probability. The financial regulators see the equivalent attributes in financial firms or pension funds (nature of the business, relationship with consumers) as an aspect of probability, along with management practices and compliance record.

The Environment Agency has three risk attributes/indicators relate to hazard or impact: these are complexity, location, and capacity for emissions. These are all inherent risks arising from the nature and location of the site itself and are mainly impact measures (amount of capacity to pollute). The probability element is the management and compliance aspects, which is assessed with respect to the site as a whole, not individual risk attributes. Other environmental regulators adopt a matrix-like approach, and assess inherent risk and management and controls with respect to each individual risk. IGAOT, for example, assesses each risk criteria on a matrix of probability and impact.

Regulators also differ as to whether the relationship between probability and impact is calculated on the basis of aggregation or multiplication. The Health and Safety Executive adopts an aggregative approach, as does the Irish EPA. In contrast, the Environment Agency is planning to move away from an aggregative approach to a multiplicative approach in

which inherent risk will be multiplied by compliance risk. It is consulting on proposals in which compliance risk will be expressed as a percentage above or below a baseline of 100%. So a better than average compliance score would multiply the aggregate of the scores for the other attributes by 95%, for example; a worse than average compliance score would multiply the score from the other risk attributes by up to 300% (EA, 2008). IGAOT also adopt a multiplicative approach, using compliance scores as the proxy for probability. In IGAOT's framework, the score for compliance history is a multiplicative criteria applied to the scores for the other six risk attributes.

An alternative to the additive or multiplicative approach is the "net risk" approach. This is used by APRA and DNB. APRA assesses the inherent risk and management and control for key risk categories and then considers the capital support available to determine the overall risk of institutional failure. The overall risk of failure is combined with the impact of failure to determine the supervisory attention index. APRA has in the past assessed management and control on a global basis across the firm as a whole. It is has now moved to a system in which each risk is measured on a "net" basis. In other words, it has started to assess the quality of management and control with respect to each risk category (*e.g.* liquidity risk, operational risk, etc.), rather than provide a global assessment of management and controls across the whole firm. This enables it to have a more granular assessment (APRA, 2008). DNB use the formula of inherent risk minus management and control risk gives net risk. The net risk figure is then multiplied by the impact figure to give a risk rating. The Financial Services Authority also assesses management and control with respect to each risk area on a net risk basis (FSA, 2006).

Weighting

A second important aspect of the design of the framework is the weighting assigned to different scores. Weighting plays a key role in all the risk-based frameworks examined, with the exception of the HSE, who have moved away from weighting scores. Weighting reveals much about a regulator's assessments of what is important, their view of risk and their risk appetite. However, it is also susceptible, like other aspects of risk-based frameworks, to "gaming" by inspectors. A number of regulators have had experience of inspectors "reverse engineering" their scores so that they obtain the risk ranking which they think is appropriate, and not that which is given by "the system". Using supervisors' less structured assessments as a general check on the accuracy of the risk model can be helpful, but reverse engineering can defeat the purpose of having the risk-based system and distort the resource allocation decisions.

Weighting can be done for a number of reasons. At base, risk attributes are weighted so that the final score enables supervisors to devote resources where those designing the framework think they will most be needed. IGAOT, for example, gives additional weight to new installations so that they become a high priority to be inspected. In the Irish EPA framework, certain activities are automatically be assigned a high enforcement category, for example incineration on land or at sea. In addition, a conviction in the last twelve months will raise the grade to one category higher than it would otherwise have been. Negative weights can also be applied to bring scores down. So where the licensed activity has not yet commenced then it is scored one category lower than it would otherwise have been (Irish EPA, 2006).

The research shows that there are other reasons for weighting. Three examples are: incentivising management; structuring supervisors' risk assessments; and structuring charges.

Weighting can be used to incentivise firms to improve their compliance. Many of the risk indicators in any framework relate to the inherent risk of the firm's activities, or, in the environmental context, its location. These are fixed, in the sense that the scores given to them will not vary between poor and well managed firms. Some regulators in the UK, notably the Environment Agency and the Health and Safety Executive, have been criticised for not rewarding well-managed firms sufficiently (NAO, 2008 overview). One approach to how this can be done is through the weighting given to internal management and compliance in calculating the risk scores. For example, in order to incentivise firms to improve their internal controls, IGAOT has deliberately assigned additional weight to firms' internal compliance so that changes in this score will have a significant impact on the score as a whole.

Weighting is also used to "correct" or structure the risk assessments of supervisors. This is particularly relevant where the assessments are qualitative. Regulators who are into the third or fourth version of their risk frameworks have progressively refined their use of weighting to take into account supervisors' behaviour in assessing risk. APRA, for example, used to give supervisors the average scores in each peer group for the different risk categories against which supervisors could compare the particular firm or fund they were assessing. However it found that supervisors were gravitating towards the peer group average in giving their scores. So instead APRA has introduced a significance weight reference points (APRA, 2008). The reference points represent the "typical" significance weights of an entity within a given peer group and are derived according to the importance of the PAIRS category to the overall business profile of the entity (APRA, 2008). The significance weight reference points are set centrally within APRA and applied to each risk category across each peer group. This enables the central risk team within APRA to ensure consistency and also to be able to calibrate the weights more easily depending on changes in the external environment. The reference points are reviewed annually or when significant events occur in the interim that would alter the risk profile of institutions within a given peer group. APRA is currently undergoing a review of the reference points with liquidity risk being given a higher weighting than in the past, for example, due to the extreme conditions in the financial markets. APRA is currently conducting further research into supervisor's behaviour to understand further what affects the supervisors' assessments of risk to see if further modifications need to be made.

Weighting is an important instrument for senior management to structure assessments being made by individual supervisors in the UK Financial Services Authority's model as well. Senior management, or the central risk team, can modify the weights assigned in the ARROW II risk model to emphasise or deemphasise the risk from certain sectors (for example sales of certain retail products), or from certain risk groups within the model (for example business risk, control risk, liquidity risk). Weighting is in turn explicitly linked to risk appetite – what level of risk the regulator decides it is prepared to accept in any one area (FSA, 2006, p. 15).

Finally, where the risk scores are linked to a charging scheme, weighting is also affected not just by risk levels but by a prior assessment by those designing the scheme of the baseline resources that are needed to supervise the organisation due to its inherent

nature. Quite simply, large, complex businesses or installations take longer to look at, so more weight is given to business complexity to ensure the charges are set at an appropriately high level.

Integrating "horizon" scanning and generic, industry wide risk assessments into the firm-specific assessment

A third key issue in the design of risk-based frameworks is the extent to which the risk assessment of individual firms or sites takes into account more generic risks arising from changes in the environment in which the firm is operating. These are particularly relevant for financial firms, whose risk profile can be significantly affected by the market environment. All the financial regulators try to identify and capture these risks, and to bring both strategic and firm-specific risks within a single risk assessment framework. TPR does this through its intelligence gathering and triage process, for example. TPR has a single data base with all the information about a fund on it. This includes fund specific information derived from returns; corporate reports; media reports; and issues in the external environment that it thinks could affect pension funding. It uses this data to derive the risk scores for the high impact funds.

Ensuring that firm specific assessments take into account these more generic risks can be difficult to achieve, however, if the regulator relies on the judgement of the supervisor alone. Both APRA and FSA have found this. The evidence as to the FSA's supervision of Northern Rock illustrated the difficulties (FSA, 2008). The answer that both have gravitated towards is again to adjust the parameters of the risk model centrally, either through adjusting weightings or pre-populating the inspectors score sheets, or both.

What to do about low risk firms – dealing with the "bulge"

For most regulators, the bulk of their regulated population fall into the low risk category. These can easily become "forgotten offenders": firms who offend but which the regulatory framework overlooks. The issue the regulator faces is what level of resources to apply to them. It obviously has to be less than it applies to high risk firms, but how low should it go? How can it identify when regulatory action is needed early enough to make interventions that could prevent the risk occurring, and how can they inform firms of the need to comply and incentivise them to do so?

Most regulators deal with this problem in one or more of three ways: information campaigns, random inspections and/or themed inspections, including sampling.

The first strategy is to use information campaigns to inform small firms of the regulatory requirements. Inspections can serve a useful function by informing firms of their obligations, particularly small and medium enterprises which typically are in the regulators' low risk categories. If inspections cease or are severely reduced for these firms, this source of information obviously disappears. To compensate, information campaigns are being increasingly used to varying degrees by many of the regulators who have risk-based frameworks. A report by the UK's NAO found that "[c]ampaigning activity plays a key role in risk-based systems of regulation in reaching low-risk businesses who might not otherwise come into contact with the regulator" (NAO, 2008f).

The HSE is at the forefront of this approach in the UK. The HSE faces significant resource constraints, and simply does not have the personnel to inspect the bulk of its regulated population on a regular basis. A firm will on average be inspected once every

14.5 years (House of Commons, 2007). In addressing this problem, it has shifted from an approach based mainly on risk, which produces a huge number of firms with similar risk profiles, to one based on achievability: what is the most effective type of intervention that it can do with respect to different types of firms, other than an inspection. It has been working on a system of "segmenting" its regulated population, in much the same way as advertisers segment their target audiences. It has been developing a number of different ways to inform and influence small and medium sized businesses in particular. In order to try to reach agricultural workers, a very difficult sector to influence, it has started going to agricultural shows, farmers' markets, and targeting information to farmers' wives. It even used the BBC radio programme, The Archers, to publicise the dangers of tractors through a storyline about a tractor fatality. To target construction workers, it is using radio and TV campaigns, celebrity endorsement, and shock campaigns. It has also co-operated with hire shops and builders merchants who have run schemes for builders to hand in old equipment and replace it with new at a substantially reduced price (financed by the shops).

However, regulators can be dissuaded from strategies of education and advice by the evaluation criteria used to audit their activities. In the food sector, EU regulations stipulate what is an accepted "official control" for the purposes of auditing food inspection authorities. These do not include offering education and advice (EC 882/2004). However, the Food Standards Agency, following research which showed the effectiveness of such strategies (Fairman and Yapp, 2005), has relaxed its own criteria for auditing local authorities to include education and advice in the intervention strategies that it will "count" in assessing their enforcement activities (Food SA, 2008).

The second strategy is to have random inspections. The reasoning is that these can be an effective way to detect some non-compliance, and if accompanied by well publicised enforcement action, can act as an effective deterrent. Moreover, as many regulators indicated, having an active enforcement policy even for low risk breaches is important as it protects the integrity of the regulatory regime. Regulatory regimes can quickly lose their credibility for regulated firms and the public if there is no monitoring or enforcement of them at all. Again, however, regulators may be restricted by their legislative and/or audit frameworks from using random inspections. The Compliance Code, for example, discourages their use, a potentially significant limitation for risk-based approaches, given the wide coverage of the Code.

The third strategy is to have themed inspections, though again regulators may be restricted from using these as the basis of rating firms. Again in the food sector, for example, partial audits or inspections (such as themed inspections) have only recently been included as one of the "official controls" that the EU will recognise as constituting inspection and enforcement activity. For others, themed inspections have been an increasingly used approach. Regulators identify particular themes or issues that they want to focus on, and inspect firm's activities in those areas alone. The choice of which firms to inspect within the theme may be random or based on a prior risk assessment.

The challenge with themed inspections is to balance attention to thematic risks with attention to firm-specific risks. The HSE moved to a topic based approach to inspections from 2002, as part of its "revitalising health and safety" approach and then its Fit 3 programme (HSE, 2004). The NAO report, conducted late in the transition, found that questions arose within the HSE as to what inspectors should do about risks that they saw during an inspection but which were not part of the "topic pack" that they were using to

assess the generic risks. As a result, there was an under-utilisation of firm specific information, resulting in the risk of making many visits to firms which fall into high risk categories for many different generic risks (although this risk was minimised by pragmatic local judgements). Moreover, inspectors felt unable to use their discretion and judgement (NAO, 2008d). Clearer communications within the HSE have since gone some way to alleviating this problem.

6.3. Risk assessment in practice

The previous section illustrated some of the key policy choices facing regulators when designing their risk-based systems, and illustrated some of the different ways in which regulators are addressing these issues. However, a risk-based framework is in practice only as good, or poor, as its implementation. All risk-based frameworks face implementation risk: the risk that they will be inadequately implemented, including the risk of "model induced myopia" – that inspectors do not look beyond the model itself. The research found that regulators faced challenges with respect to three main aspects of implementation: collecting and managing the data in order to identify and assess risks; the performance of the risk assessment process, and the design and operation of the internal systems of governance over the risk-based approach within the regulatory organisation itself.

Collecting and managing data

Data is critical to the design and operation of risk-based frameworks, and can pose a significant problem. Many of the regulators examined here, and evidence from other reports on risk-based systems, emphasise the difficulties that arise because of data (*e.g.* DNB, 2006, p. 56; IOPS, 2007; NAO, 2008f).

Regulators usually have too little of the information they need, and too much of the information that they do not. If they have too little data, they obviously need to collect more. However data is highly resource intensive to collect both from firms and from elsewhere. As we have seen, risk-based systems usually incorporate information about matters outside the individual firm, for example on the geology, flora and fauna and social geography of the location; or on the conditions in the markets or particular economic sectors. Even if this and other relevant data is held somewhere in government, it is often dispersed across different governmental bodies or between central and local government officials. This can pose problems of co-ordination and delay. For example, the UK OFT has found that in developing its ability to target higher risk activities such as mass-marketed scams, it has to co-ordinate with local trading standards officers. However the lack of an integrated management system for sharing intelligence, the uncertain status of the OFT as leader of the project, and difficulties in funding a regional intelligence network have all posed obstacles (NAO, 2008e).

There is a significant difference between regulators who operate through a licensing regime, such as environmental and financial regulators, and those who do not, such as food and occupational health and safety regulators examined here. Those who operate a licensing or even notification regime have at least some way of knowing who their regulated population is and through the licensing process they have a means of obtaining information from those firms (although difficulties remain in identifying those who operate illegally without a licence, and their information gathering powers may be truncated even with respect to licensees). Most licensing regulators use the introduction of a risk-based approach as an opportunity to reform their licensing process in order to get

the information it needs about its regulated population. Some are also using the new data requirements to filter out the industry. In the UK, the OFT has enhanced the data requirements of firms for consumer credit licences to require them to give sufficient information to demonstrate competence and the adequacy of their internal management. It has found that many small operators who do not really need a licence do not want to go through that process, and so they are not applying. Others are being forced to think about their business in a different way.

In contrast, those regulators who regulate across industries in regulatory regimes where industries do not require a licence face particular problems in getting information as there is no licence process to alert them to who is doing what. The UK HSE, for example, is aware of its own tendency just to focus on the largest firms as these the most visible, although it is taking steps to use a wider range of information sources (NAO, 2008f).

On the other hand, it is easy for regulators to be swamped with data, and as a result to become locked in an endless task of processing rather than evaluating the information that comes in. The UK Pension Regulator's (TPR's) predecessor, the Occupational Pensions Regulatory Agency, for example, received 56 000 notifications in 5 years. TPR still gets around 2 500 notifications and queries per month, ranging from notifications of trustee details to information on major corporate transactions. It has a two level filtering system to prioritise them. Customer support deals with the most straightforward inquiries and notifications of minor breaches. More serious issues are sent to "triage" for analysis. Triage usually reduces the number down to about 100 high risk issues which then become cases. Cases are then directed to specialist practice teams depending on the issues they raise: corporate risk governance; scheme specific funding; and pensions administration and governance (NAO, 2007; TPR, 2006, 2008).

As a result of problems in getting the right type of data initially, regulators often design their initial risk-based systems on the basis of the information they have already got or can easily and quickly acquire, rather than on the basis of the information that they need. Indeed, regulators may explicitly design the first version of their risk-based system in such as way to generate as much data as possible, with a view to refining the framework further once it has sufficient information on which to make a more informed risk assessment. Later versions of the risk-based regime can then reduce the information requirements for low risk firms once the regulator has sufficient data to identify them.

For all regulators using risk-based systems, the IT system is a critical instrument for data management. The IT system which processes the inputs of the risk-based framework is used to collate data and to organise it. Nevertheless, a common criticism of regulators, including those with risk-based frameworks, is that they fail to make full use of the information that they have (*e.g.* NAO, 2008f). This links in part to the question of how to integrate "horizon scanning" or broadly contextual information into firm specific risk assessments. Moreover, knowing what information to seek, and managing it, is critical to knowing what new risks may be relevant, and thus to the continual modification of the risk-based framework.

Performing risk assessments

In talking about assessment, it is important to distinguish between the collection of information with regard to the risk indicators, and then the assessment at to what risk category should be assigned based on that information. Two key differences between

risk-based frameworks is who collects the information for the risk indicators, and the extent to which individual supervisors or others have discretion in determining which risk categorisation should be applied given the information gathered. Risk assessments are inherently judgemental processes. Regulators vary in the extent to which they try to "design judgement out" of their frameworks, or, if they cannot design it out, structure how it is used. There is also a close relationship between the allocation of responsibility for gathering information for each risk indicator and the degree to which the risk scores are automatically assigned.

There are four different ways in which information for the risk indicators is collected: by the firm, a contracted third party, a municipal or state government, or the regulator itself.

In environmental regulation in the UK and Ireland, the firm itself provides the information with respect to each risk indicator. In environmental regulation in Portugal, it is intended that a contracted third party obtain the information. In food safety regulation in the UK and waste management regulation in Ireland, local governments perform the assessment. For all the financial regulators, the regulator's own supervisors perform the assessments.

In the environmental frameworks in Ireland, the UK and Portugal, the indicators are objective measurements or "yes/no" answers, for example, does the firm have a management system which is externally accredited. A risk score is assigned in the framework to each measurement (*e.g.* > 10 tpa is low risk, 10-25 tpa is medium risk; < 25 tpa is high risk). There is thus very little scope for judgement in assigning the risk score. Judgement, of course, is exercised by those designing the framework, for example to determine whether emissions over 25 tpa should be high risk, or whether that figure should be higher or lower. But at the level of making individual assessments, judgement is designed out of the assessment process as much as possible.

This design is deliberately to enable the firm to complete the assessment and to ensure consistency of responses. It does not, of course, ensure accuracy of responses. The Environment Agency validates the responses through inspections. Baseline inspections for those subject to the individualised risk assessment process occur annually (recall that only the highest risk installations are subject to the bespoke risk assessment); for those in the higher risk groups, they occur more frequently as determined by the Opra score. The Irish EPA validates the responses from the operators through a desk-based assessment of the returns submitted. The Portuguese environmental regulator, IGAOT, in contrast, will contract out the task of completing the risk indicator forms to third parties, to ensure from the outset that the information it has is valid.

The rationale for self-completion by firms of the risk indicators form is based on pragmatic and strategic considerations. Pragmatically, firms have the information and so it makes sense that they should complete the forms. Strategically, the regulators argue that the process of completing the forms means that the firms start to recognise their own risks, to see their operations from the regulators' point of view, facilitates "buy-in" from the industry, and reduces the potential for disputes over the categorisation. Even if the regulator raises the categorisation when they verify it, experience suggests that the number of disputes is reduced.

The collection of information for the risk indicators and the assignment of risk scores are more closely combined in the risk frameworks of the financial regulators, and this two-stage process is often collapsed into one. In many areas of assessment, the range of variables is so great that the framework cannot envisage all of them and assign a risk score in advance

to each, or at least not without becoming overwhelmingly complex. Many of the assessments are therefore subjective, and not based wholly on quantitative measures or yes/no answers. The translation from the information that the supervisor collects into a risk score is thus a matter of judgement. This type of risk-based framework poses quite different issues. Self-completion by firms would be a more significant step, for regulators would not only be relying on firms to give accurate responses, but to give responses which involve qualitative assessments. The arguments for self-assessment may still apply, but the level of judgement involves means that regulators are likely to be less comfortable with self assessments without far more extensive validation than regulators in the environmental sector requires, given the scope for inconsistency in assessments that would arise.

Finally, the risk assessments may be performed by local authorities or municipal governments under the guidance or direction of a central state regulator. This is the model used in the UK's food safety regulatory regime. The UK Food Standards Agency's Code of Practice sets out the risk framework, the minimum levels of inspections for each risk category, and the parameters of the compliance or interventions policy. It has no powers however to determine what level of resources that local authorities should spend on food inspections, though it does have powers to take over their responsibilities if it considers that they are being inadequately performed. `

Internal governance of the risk-based system

There is a close relationship between the organisation of the risk assessment process, in particular the degree to which completing the assessment relies on individual discretion, and the organisational structures for governance of the risk-based framework within the regulator. In the environmental regulators examined, inspectors have limited discretion for assigning the risk score (though as we will see below, they still have discretion as to how to respond to individual risk scores). These regulators need a process to validate individual, firm level information and to review and periodically recalibrate the risk framework. However there is therefore less need for an internal governance process to ensure consistency or accuracy of their judgements. There is a need to ensure that when inspectors validate firms' own assessments that they look at the appropriate things and have the technical ability to assess the firm, but the main challenge of consistency comes with respect to supervisory response, not the risk assessment *per se*.

In contrast, in those risk-based frameworks in which supervisors have a considerable degree of discretion in assigning risk scores, regulators have to ensure that supervisors are consistent and accurate in the scores that they give. For these regulators, internal risk governance processes are central, and the introduction of a risk-based framework often entails wide-ranging and on-going changes in the regulators' internal organisational structures. Both APRA and FSA, for example, are on their third or fourth model of internal governance.

Risk-based frameworks that are based on supervisory judgement have three main challenges: how to ensure the quality of supervisors' assessments; how to ensure consistency; and finding the balance between central control and supervisory discretion in assigning risk scores.

Quality of assessments

There are three key issues with respect to quality: training; integrating contextual risk analysis into firm level risk analysis; and understanding supervisors' behaviour in making their judgements.

All regulators who have risk-based systems emphasise the need for training. However, training has to be not just in mechanics of the assessment, but in the whole philosophy of risk-based regulation. The most common mistake that early adopters of risk-based frameworks said that they had made was that they assumed that supervisors would know what a risk-based assessment was, and that therefore they simply needed to be trained in the IT, in how to fill in the assessment spreadsheets. What they found was that supervisors were not really aware of the distinction between a compliance based approach and a risk-based approach to supervision. This problem is not confined to risk-based frameworks which are based on supervisory judgements, and environmental regulators reported the same problem. In those with frameworks based on structured risk classifications, *e.g.* the environmental and food regulators, this issue it manifested itself at the stage of deciding what enforcement action to take; in supervisory judgement frameworks, it manifested itself at the level of assessment as well.

Regulators have different expectations as to the training that their inspectors or supervisors have to undertake. In the health and safety context, for example, the HSE requires inspectors to undergo a two year training programme, take a formal qualification, have ongoing assessments by the peer group and specialist training inspectors. The HSE also provides extensive internal guidance on the objectives and rationales of the inspection with respect to each topic, with examples of the types of responses and interventions they should make, and what is good and bad practice for inspectors. In addition, its Enforcement Management Module provides guidance on risk assessments and on the appropriate responses inspectors should make.

The second issue is integrating contextual risk analysis into the firm level risk analysis. As noted above, market context can have significant effects on the risk profile of individual firms in the financial sector in particular. All the financial regulators examined here have specialists responsible for performing this analysis, usually in a specific division within the regulator. However, as the FSA's experience of supervising Northern Rock illustrated, it can be difficult to ensure that supervisors integrate that risk analysis into their firm-level assessments (FSA, 2008).

The third issue is understanding supervisors' own behaviour in performing the risk analysis. Risk assessments are inherently judgemental, but are critical to the regulators' understanding of its regulated population and to how it responds. Regulators therefore need to understand how supervisors behave when making those judgements. Regulators who are into their second or third generation of risk-based frameworks are developing an awareness of how they need to structure the assessments to adjust for supervisors' behaviour.

APRA, for example, used to give supervisors the average scores in an industry for the different risk categories; however it found that supervisors were gravitating towards the industry average. So instead APRA moved to a significance weight score. APRA has always weighted the different capital support categories; it now weights the different PAIRS categories. Significance weights are derived according to the importance of the PAIRS category to the overall business profile of the entity (APRA, 2008). The significance weight score is set centrally within APRA and is applied to each net risk category. This enables the central risk team within APRA to ensure consistency and also to be able to calibrate the risk scores more easily depending on changes in the external environment. Liquidity is currently being given a much higher weighting than in the past, for example, due to the

extreme conditions in the financial markets. APRA is currently conducting further research into supervisor's behaviour to understand further what affects the supervisors' assessments of risk to see if further modifications need to be made.

Others have also begun to incorporate an understanding of how supervisors assess risk in its risk model. Through its validation processes one regulator discovered that supervisors would over-estimate the quality of management and controls to a relatively high degree, around 30%, and moreover that this over-estimation was consistent across supervisors. Helped by the consistency of the judgements, the regulator is able to adjust the basis of the calculations of the risk scores to take this over-estimation into account.

Further, some in some areas is can be difficult to identify the difference between a risk and a control. The financial regulators are finding this, perhaps particularly at this time: that supervisors may assess certain features of the firm's risk management strategy to be controls, the risk division see them as risks. For example, the structure of control systems can themselves be risks if they structure incentives in a particular way or if they cannot counteract the incentives structured by the systems for awarding pay and bonuses

Current events in the market raise fundamental questions as to when a control becomes itself a risk, and indeed the moral hazard created by the control structure itself.

Consistency

Consistency is closely associated with quality. All regulators with these frameworks found that the internal governance structures were a key issue in ensuring consistency of assessments across a large number of supervisors, and that it was not easy to get these right. In addition to training, key issues were ensuring that internal comparisons and validations were made of supervisors' assessments.

Again, regulators have experimented with different structures. APRA began with PAIRS panels. These were panels of senior management, and they would go through two or three risk assessments in depth with the supervisors, challenging them to ensure accuracy and consistency in assessments across the organisation. However, experience showed this was a relatively cumbersome process in practice, and so APRA has moved to PAIRS forums. This is a more group wide approach to the benchmarking process. The forum is comprised of senior management, other supervisors and the appropriate risk specialist. Around 10 entities are randomly picked from each group of institutions and considered. The forum discusses with the supervisors how they arrived at their scores in order to check for outliers and discus the criteria that supervisors are assessing against. The forum does not have the power to change the rating; APRA considers it important that the final decision lies with the supervisors, though supervisors are likely to change the score if it has been successfully challenged in the forum.

Issues of consistency vary with the number of supervisors involved. The UK Pensions Regulator has only 20-30 people performing assessments, and does not conduct inspections; its problem instead is filtering information that comes in. It has "triage" system and Tasking Co-ordinator Group meeting which decides whether to intervene in a particular case. If it does decide to intervene then set up a taskforce with a case manager, lawyer, actuary, and sector specialist business analyst. Criteria for intervention are set out in TPR's "business rules". These determine how certain types of information are dealt with. An example is the business rule on scheme recovery. Pension schemes that are in deficit must submit a recovery plan to TPR. The plan details how the deficit will be recovered and

over what time period. TPR has created a set of trigger points that indicate when further action should be taken. The business rules are used by the staff to guide their analysis of the deficit recovery plan (NAO, 2007).

A key issue is how to ensure consistency of risk assessments and an "all round view" of risks without creating overly cumbersome committee/panel structures and paralysing the organisation in procedures. Some find that the obstacles to getting information in on all the different risks from a wide number of inspectors or inspectorates, each of which is looking at a particular part, is simply so challenging that it is rarely done. For those that do try to establish a system wide view as part of their standard operations, it is easy for internal structures to proliferate. One regulator reported that the internal assessment system at one point consisted of fourteen committees at four different levels. This clearly affects the speed and responsiveness of the regulator, something which is particularly relevant where external market conditions are highly relevant for risk assessment and where these changing rapidly. It is hard to have a "real time" risk analysis if everyone in the organisation has to have a view. As one regulator said, the central risk unit could do the evaluation but that would not be seen as valid, as it had not been validated by all the different units within the regulator. There is thus a tradeoff between ensuring accuracy, consistency, and "buy in" from across the regulator with speed and responsiveness.

Balance between central control and supervisory discretion

Within all the regulators, there is a separate set of officials responsible for the design and ongoing maintenance of the risk-based system. This unit evaluates the framework, and sets the risk parameters on which the gradings are based. The relationship of this unit with the rest of the regulatory organisation varies. It may be focused specifically on risk analysis, or have a wider role. APRA, for example, recently established a Supervisory Framework team, which is a single team across APRA dealing with all the different industries, and which is responsible not only for the maintenance and development of the risk framework, but monitoring supervisory activity across the whole of APRA, training supervisors and producing guidance for them.

One of the issues that regulators have found is how to balance control by the centre over the risk assessments with local discretion. The degree of central control exercised over the risk assessments of supervisors varies considerably (see also IOPS, 2007).

Risk-based systems, as we have seen, can potentially place heavy reliance on the exercise of discretion and judgement by supervisors and inspectors to ensure that risks have been properly identified, to assess them and to assess whether the preventive measures taken are adequate to control the risk. On the other hand, they remove the discretion of who to visit and when; and perhaps what to look for. Inspectors can then feel devalued. For those regulators who use the self-assessment process, inspectors lose their role to categorise firms. So an inspector might have thought that X was a "good company" but it comes out as high risk, and so the inspectors' assessment is displaced. That kind of personal judgement gets removed from the assessment. This is a significant shift in practice and culture. When risk-assessment frameworks were just being introduced, many inspectors found this hard to accept. Ultimately the central risk teams have found they have to allow inspectors to make representations against a risk score, but in practice few categorisations have been changed because, if they differ from an inspectors' own past experience of a firm, they need a very good reason to have it changed.

Much depends on the internal culture within the organisation. In some regulators, contrary to the example above, the inspector or supervisor can be seen as "king" within the organisation, and as knowing the firm better than anyone else. This can make it very hard for central risk unit to get organisation to move to a "portfolio" approach rather than one led by individual risk assessments, or indeed to get supervisors to change their assessments. It can make for internal battles, as it is hard for supervisors to accept that "their" firms are not as significant for the regulatory organisation, and thus as deserving as resources, as someone else's.

Some regulators allow senior management in different areas to customise the model and adjust the weightings and aggregations of risk scores in their industry areas. Regulators have found that this has helped to engage managers; as one member of a risk team commented, "they can play with it". However it had the effect that the risk scores went up, as everyone thinks their area is more risky than anyone else's. Central risk units then find themselves having to "rebase" the scores to scale them down, and readjust them between divisions in line with its own evaluations to ensure that resource allocation was not distorted.

One technique used by a number of regulators is to "pre-populate" the risk scores. In environmental risk-based systems in the UK, Ireland and Portugal, for example, all the scores are automatically assigned by the framework. In those systems which rely more on supervisory judgement, pre-population has also developed as a technique to ensure some central control over risk weightings. For those who were the "early adopters" of risk-based systems, pre-population developed over time, and so now tends to be characteristic of a second or third generation risk model. Those introducing them now and learning from this experience have benefited from this learning to introduce the technique straightaway. Pre-population can be an extremely useful way in which the centre can structure the judgement of supervisors. Indeed, some financial regulators have found that the only way to ensure that supervisors capture the external risks which it sees as relevant to a firm, for example, is to pre-populate the risk scores.

What do regulators use their risk-based frameworks for?

Allocating resources

One of the purposes of a risk-based framework is to facilitate the efficient and effective allocation of resources. Its presumed role in achieving this purpose is the reason why the UK central government is requiring all regulators to adopt risk-based systems. It is also the reason why the European Commission is incentivising regulators to adopt it in the environmental sector.

Three main questions arise: to what extent do resources follow risks in practice; do regulators in fact have the resources to inspect all the firms that score as "high risk" on their risk scorecards, and what other uses do regulators make of their risk-based frameworks?

Mobility of resources – do they shift and are they adequate to cover all the "high risk" firms? One of the main reasons that risk-based systems are advocated as part of the "better regulation" drive is that they are meant to be a tool for efficient resource allocation. Regulators agree that broadly speaking resources do shift between the main risk categories, but that it is harder to get resources to shift in lines with more fine grained changes in risk assessments.

As explained above, regulators frequently divide their regulated population on the basis of broad impact measures. These do broadly determine resource allocation. So in environmental regulation in England and Wales, a Tier 1 firm receives a tiny proportion of the attention of a Tier 3 firm, for example. This categorisation often determines when the bespoke risk assessments that we are discussing here are done. Difficulties arise in determining which firms within this category, or possibly two highest categories (depending on how many categories there are) require the most resources. Many regulators find that in practice it is hard to ensure that resources shift in accordance with the risk assessments within these higher risk bands, for a number of reasons.

First, risk-based regulation means not doing something; it is hard for regulators to decide what not to do. Once the lowest risk firms have been discounted, as it were, it is difficult for all risks within the higher risk bands not to seem equally as important. Moreover, it is hard for the organisation as a whole to adopt a "portfolio" approach to managing its most significant risks, and to see beyond an individual firm, or firms in a particular sector. One regulator has introduced a two stage process to determining resource allocation with its very senior management and Board. It asks those at the top of the organisation to set a particular *quantum* of risk and resource allocation – to adopt a baseline of say 100, and then rank firms above or below that baseline. But as one pointed out, "it's a zero sum game, and [top management] find that hard to understand, that if we put resources here that means they're not available somewhere else".

Second, there may be reasons for resources not to be determined by the score in particular instances. The Environment Agency is clear that a firm's Opra score is a guide to resource allocation, but only that. Ultimately decisions on resource allocation are made at the regional level, and various factors can modify the resource allocation decisions suggested by the Opra score. So a site will have higher priority than the Opra score would suggest if, for example, the installation or site has been given a lot of improvement conditions, if it is new to the sector, and if it is a contentious site, one that gives rise to a significant number of local objections, a point discussed further below.

Third, there simply may not be enough resources to monitor all the high risk firms in a way that the system might envisage. This may be because the risk scoring is not sufficiently fine grained, but it may also be that there the regulator is simply under resourced. As one regulator commented: "It's very hard to match complexity of the legislation to the risks and then to capacity – it's not one on one… we have more high risks than we have capacity". Those who are just introducing their risk-based systems recognise that they will have to feel their way, to some extent, on the issue of resource allocation. IGAOT, for example, intends to inspect its high risk sites annually; its medium risk ones every two years, and to use random inspections for its low risk firms unless there have been complaints. But it recognises it will have to see what the scores come out as before it can make a final decision. The Irish EPA is in a similar situation. In practice, there may be too many high risk firms for either regulator to perform their desired level of inspections given their current resources.

The HSE in particular has found that in recent years it has had to divert an increasing proportion of its inspectors' time to investigating accidents, rather than performing preventative work such as inspections (NAO, 2008). It is under a mandatory obligation to investigate all major injuries and fatalities. The time taken by this activity has increased partly because of the complexity of the issues, partly because of greater concern by families

of the deceased, and partly because firms are more likely to challenge formal enforcement action than they were in the past. The result has been that fewer resources are available for inspection, even of the higher risk businesses, than there have been previously.

It is hard to know whether this is a problem, however. As many regulators observed, it is difficult to establish what the right number of inspections is for the regulator to be able to say with any confidence what the level of compliance is in a particular area of activity, and to be able to improve it.

Moreover, inspections are performed to achieve a number of objectives: to meet legal requirements; to identify breaches and apply sanctions; to monitor compliance levels and target problem areas; to help businesses comply with the regulations; to prevent major incidents and (critically) to maintain confidence of stakeholders (*e.g.* NAO report, p. 17). Risk-based approaches conflict with the achievement of some of these goals. In particular, helping businesses comply and maintaining confidence of stakeholders require higher levels of inspection that risk-based systems would normally allocate. Yet these goals still need to be achieved. A key issue and point on which regulators often differ between themselves, and with politicians and other stakeholders, is the extent to which inspections should continue to play a valuable role in their attainment.

What ultimately drives resource allocation, however, is the political context and the risk to the regulators' own reputation. As noted above, some regulators routinely factor in public perceptions and the risk of damage to their own reputation in allocating their inspection resources, others do so implicitly. The UK Food Standards Agency, the HSE and the Environment Agency deliberately take into account public perceptions in allocating inspection resources and believe they would be heavily criticised if they cut back inspection activity. This has a significant bearing on the allocation of their resources. The HSE and Environment Agency believe that after their preventative work, the public expectation is that they will investigate and prosecute companies in the wake of accidents or pollution incidents. As noted above, HSE spends over half its front line regulatory resources on accident investigations (NAO overview, p. 17). The UK Pensions Regulator clearly states that firms in the intensive monitoring are those that pose highest risk to objectives, risk is also defined as "risk being that we may be perceived as not making a difference" (TPR, 2006, p. 50).

There are some firms or risks that in political terms a regulator simply cannot leave alone, regardless of the probability. As one commented, "events force you up the probability curve". The higher the political salience, the lower the probability level at which the regulator will intervene. Political risk here is critical in determining a regulators' risk appetite and its risk tolerance, and thus the allocation of regulatory resources; regardless of what the impact and probability studies would otherwise say.

Other uses of risk-based frameworks. Allocation of inspection resources is only one use of a risk-based framework. Regulators also use the frameworks for a number of immediate purposes, as well as to achieve the broader motivations indicated above. Principal other uses are:

- to set fees and charges;
- to provide information for reporting purposes, particularly in the environmental context;
- to gather information on the regulated population; and
- as part of broader strategy and objective of improving management engagement and compliance performance.

The use of the framework to set charges is common amongst the environmental regulators. However, this does mean that much of the risk grading is attributable to fixed attributes of the site, notably its scale and complexity. This is because complex sites take longer to inspect, and so consume more inspection resources. The risk score therefore has to be high for such sites to enable the charges to be recouped. The extent to which the charging structure drives or influences the framework does depend on whether charges are applied on a cost-recovery basis or not. The Environment Agency has to apply charges on this basis, and this has raised a number of issues which potentially cut across the pursuit of a risk-based approach. In particular, inspectors feel that they have to spend longer on such firms as those firms have paid more (NAO, 2008a). The annual enforcement charges assigned by the Irish EPA to the operators take into account the enforcement category arrived at through completion and validation of the methodology. In general, the higher the enforcement/risk category, the greater the annual enforcement charge which the operator has to pay the Irish EPA.

The frameworks are also used by many of the regulators as part of a broader strategy to engage management. This is rationale is particularly evident in those frameworks using self assessment, such as the environmental frameworks and the UK Office of Fair Trading's new approach to licensing. Through the self assessment process the regulator is attempting to ensure firms engage with the regulation and moreover see their operations from the regulators' perspective. It can also be a way of handling the inheritance of a regulator from a previous regulatory regime. In the case of the UK Office of Fair Trading, the previous routine approach to licensing for consumer credit meant that firms simply applied for as broad a license as possible, and the regulator had little idea of what areas they in fact were operating in. The OFT thus has 156 000 current licence holders. Asking them for information in detail in the application process is a relatively efficient way of getting information on their business (subject to validation) and prompting them to reduce the number of different types of consumer credit business for which they apply for a licence.

Performing inspections in risk-based frameworks

One of the most significant challenges for regulators moving to risk-based systems is changing the culture and skills of inspectors. All regulators examined whose risk-based systems have been running for some years have found that it takes at least two years for inspectors to move towards a risk-based approach to inspection. And as the FSA's experience with the supervision of Northern Rock illustrates, it can take far longer.

Four key issues emerged from the research with respect to inspections: the training and re-skilling of inspectors; how to avoid false positives; how to balance a focus on outcomes with a focus on compliance; and how to manage risk-based inspection systems in a federated inspection structure.

Training and re-skilling of inspectors

Risk-based frameworks have significant implications for inspectors and the inspection function. The shift to a risk-based approach often requires a fundamental change in culture, a different analytical approach, a different understanding of the role of inspectors and supervisory staff, and a new skill set. All the regulators examined here, and those examined by others, have found that this is a key challenge in introducing a risk-based system (IPOS, 2007; NAO, 2008f, p. 17).

A shift to risk-based inspections is particularly challenging where the organisations involved in inspection previously had a long practice of routine processing of information or routine inspection processes. These changes can prompt hostility to the risk-based regime from some inspectors. As noted above, by its very nature, risk-based frameworks significantly curtail the scope for inspectors' discretion in determining how to plan inspections, who to inspect, and what to inspect for. It can also be difficult for inspectors to accept that they no longer need to spend too long on particular firms, as it calls into question the validity and usefulness of the way they have performed their roles previously.

Often regulators find that in order to begin to change the inspection culture there has to be a shake out of the current supervisory staff, and new people hired or brought in on secondment. However, even in those regulators who have operated a risk-based framework for some years, firms complain that inspectors are insufficiently skilled and knowledgeable to make risk-based judgements, and that they still have a "tick box" mindset (NAO, 2008a, p. 31).

Many of the regulators who have had risk-based frameworks for some time admit that in hindsight they spent too little time on training inspectors, and/or that the training they did was focused on the wrong things. Frequently, training was given on the IT system and on how to fill in the risk assessment forms. However, what was neglected at first was training in the whole philosophy of risk-based regulation. As one regulator commented, "we thought they would just get it, just understand what risk-based meant, but they didn't".

Avoiding false positives and false negatives

One of the problems that regulators with some years' experience of risk-based frameworks have found is that the system can return false positives or negatives, depending on how it is designed. Where a supervisor or inspector is not sure of how to grade a particular risk, in some systems they can leave this blank. If the IT system underlying the framework automatically defaults to a low risk score, the result can be a lot of false positives. It may be that the score was left blank because it was low risk, but it may also have been left blank because the supervisor or inspector did not look at the issue or did not understand it.

Regulators have met this problem in different ways. The Dutch and Portuguese environment regulators' frameworks cannot be left blank, so one of the appointed solutions is to fill a medium score to those criteria for which there is no available information, to avoid giving weight on high or low priority which could lead to false. The UK Financial Services Authority's revised framework, Arrow 2, requires supervisors to enter a judgement to avoid leaving "dark holes" where the risk score does not properly reflect risk because of an under estimation by supervisors or because it is simply out of date, though as Northern Rock illustrated these "dark holes" still exist.

None of the regulators examined has a system in which the person doing the assessment is required to state their confidence level, however. In contrast, the peer review process for some research councils requires referees to state how expert they are in the particular research area and how confident they are about the rating they give. Some regulators are thinking of introducing such a system, although there are issues as to whether inspectors or supervisors will in fact admit to lack of confidence.

How to balance focus on outcomes with focus on compliance

The shift from a compliance approach to a risk-based approach can also be problematic because of the legal framework in which regulators have to operate. Regulators are often charged with implementing an existing set of legal requirements which are not outcome focused, and which they are unable to change. It may well be that breach of a particular requirement does not affect the risk or outcome. The fact that there is a disparity suggests the rule should be re-written if not removed, but often it is not within the regulator's power to make these changes. In the EU context, it can often require a change in EU legislation. Leaving a number of breaches unsanctioned can reduce the credibility of the regulatory regime as a whole, however. For this reason, inspectors can resist the move to a more risk-based approach.

Federated inspection systems

Federated inspection systems pose particular problems. The extent to which the central or federal regulator, or regulator operating at the level of central government in non-federal systems, can influence what happens at a local level varies with the constitutional and political context of each country. Co-ordination problems are clearly enhanced where the central regulator can exert little control. However, even in systems where the central regulator does have powers over the inspection processes of local or regional authorities, there can be problems with the co-ordination of inspections and consistency in risk assessments. For example, in the UK, in food regulation there have been problems of "join up" between local authorities and the central agency. The Agency sets its own priorities for food safety, but as these are not legal obligations, they are not reflected in the Code of Practice. The result is that inspections and the regulatory priorities are not integrated. Similar problems can arise across regulators with a large number of regionally dispersed inspectors (see *e.g.* NAO, 2008d).

Compliance/enforcement policies and risk-based frameworks

How closely the regulators' risk assessments are linked into a particular enforcement approach is a significant point of variation between the different risk-based frameworks. Many regulators have enforcement policies or compliance strategies. These may categorise firms on the basis of their attitude to compliance, as in the case of the Environment Agency and HSE, for example. The enforcement strategies may themselves be risk-based in that they incorporate an assessment of the likelihood of success of formal enforcement action, such as that of VROM or the Financial Services Authority. Often, however, there is no direct link between the risk category of a firm and the enforcement strategy that the regulator will adopt. Notable exceptions are APRA, DNB, and VROM.

APRA, whose model was followed by DNB, integrates its risk assessment framework to the type of supervisory response it will take. It uses PAIRS to determine a firm's risk level. PAIRS is integrated with SOARS – the Supervisory Oversight Assessment Framework. The development of both PAIRS and SOARS was shaped by the failure of the insurance firm, HIH. This event had revealed both the weaknesses in APRA's existing risk-based frameworks for assessing financial institutions and the absence of an effective culture or practices of supervision and intervention. SOARS was devised to address that failure, and is deliberately intended to create a more pre-emptive and effective supervisory intervention culture within APRA, and to improve consistency in its supervisory interventions (Black, 2006).

SOARS has two components: a supervisory attention index and a supervisory stance. The supervisory attention index computed as the geometric average of the probability (risk) index and the impact index. Although supervisors have discretion as to exactly which intervention and enforcement tools they use, the SOARS index sets the amount of supervisory resources each institution is likely to require, and the supervisory stance that is to be adopted in terms of its relative intrusiveness, intensity and directiveness.

Table 6.5. **The SOARS grid**[1]

		PAIRS Probability rating				
		Low	Medium		High	Extreme
PAIRS impact rating	Extreme	Normal	Oversight	Mandated improvement	Restructure	Restructure
	High	Normal	Oversight	Oversight	Mandated improvement	Restructure
	Medium	Normal	Normal	Oversight	Mandated improvement	Restructure
	Low	Normal	Normal	Oversight	Mandated improvement	Restructure

1. The grid is published widely in APRA documents. See for example, APRA's Risk Rating of Superannuation Funds (Insight, May 2004); APRA, *Annual Report*, 2003.

The intervention settings for the supervisory attention index and the supervisory stance are set by APRA's senior executive and Members. They are currently torqued towards earlier and more interventionist action for larger firms, again a direct consequence of HIH.

In its initial form, APRA's SOARS framework set the level of supervisory resources and supervisory approach, but left the choice of individual intervention plans to the supervisors. APRA has also begun to give the same attention to the supervisory responses adopted with respect to individual firms within each category as it has to risk assessments. The discussions of the PAIRS forums have begun to integrate discussion of the risk assessments with discussions of what the supervisory response should be. APRA is also establishing SOARS panels to establish the same level of scrutiny over the supervisory approach being adopted as they currently have over the risk assessment.

VROM also integrates risk assessment with supervisory response. It has integrated inspection and enforcement teams, which includes members who are specialists on the effectiveness of different intervention strategies. The intervention strategy is linked to the level of risk (VROM, 2004). When an organisation is ranked as red, which is high risk, then a more severe approach is taken. They have intervention specialists and members of the prosecution authorities within the project teams (soil, water, air quality and safety), who work with the inspectors and other team members to explore what would be the best type of intervention to make. VROM have a well articulated Compliance Strategy. This seeks to combine a "task-oriented track", which focuses on the rules to be enforced, and a "problem orientated track", which focuses on the problems to be addressed. High priority is given in enforcement to breaches of rules which pose a high risk and with respect to which there is a high-non-compliance rate. Medium priority is given to areas with respect to which there is low non-compliance but which are high risk, and to which there is high non-compliance but they are low risk. Using the media to draw attention to issues and non-compliance is a key part of the enforcement strategy for medium risk occurrences. Low priority is given to enforcing a rule if there is low non-compliance and it is low risk. The form that the enforcement action takes then varies with an assessment of how enforceable the rule is, the firm's motivation for non-compliance, how it is likely to respond to intervention, and whether a broader approach to tackling the problem is required. VROM uses the "table of 11" used by the Dutch Ministry of Justice as a framework for determining what intervention to take.

Table 6.6. **Dutch Table of 11**

Aspects of spontaneous compliance	1. Knowledge of the regulations. 2. Cost/benefit ratio. 3. Degree of acceptance. 4. Loyalty and obedience. 5. Informal monitoring.
Aspects of monitoring	6. Informal report probability. 7. Monitoring probability. 8. Detection probability.
Aspects of sanctions	10. Choice of sanctions. 11. Severity of sanctions.

Communication of results

Communication both of the nature of the framework and of the results of the risk assessments poses a number of issues. With regards to the framework, the Irish Environmental Protection Agency, for example, found that describing sites in terms of "risk" caused too much confusion, so it deliberately named its framework an environmental "assessment". With respect to communication of results, confusion can often arise as many regulators have found that firms do not necessarily understand the results of the risk assessment or the implications for their relationship with the regulator. Regulators have found that they need to pay greater attention to this aspect of communication than they at first thought.

Regulators also adopt quite different approaches to whether they communicate individual firm's risk assessments to the public or not. The financial regulators do not publish risk assessments, largely out of concern that they will be misunderstood by the public and damage market confidence. In contrast, the Environment Agency does publish the Opra risk assessments of installations via its websites and through its Spotlight reports.

The issue of whether and how to publish the outcomes of the risk assessments has come into sharp relief in the context of UK food safety regulation. Following the example of the Dutch authorities, many local authorities in the UK have started publishing "scores on the doors" of the food establishments that they inspect. There are now over separate 200 schemes run by local authorities, many of which uses a different scoring system. Many of these incorporate all or aspects of the risk score derived from the inspection process. The Food Standards Agency is currently consulting on developing a single, nationwide framework for "scores on the doors" (Food SA, 2008b). Publishing "scores on the doors" can be a very effective way of harnessing consumer power to reinforce the regulatory process. However, there are concerns that the "scores" can only give a snapshot picture of the state of the establishment at the time of inspection, and moreover that it would not be appropriate to incorporate all aspects of the risk score into the "score on the door" as the two have quite different purposes. The principal argument put forward by regulators in a range of sectors for not publishing scores is that they will be misinterpreted. Risk assessments are internal tools used by regulators for a number of purposes; they are not assessments of the quality or even compliance levels of the firm itself.

6.4. Evaluation of risk-based frameworks

Evaluating the effectiveness of regulation is a significant challenge. In order for their risk-based approaches to be effective, regulators have to know whether they are in fact applying the right level of resources to the right issues. But as noted above, it is hard to

know how many inspections to do when the impact of each one is hard to evaluate (see also NAO, 2008f, p. 17).

Traditionally regulators, and their auditors, have been very good at counting what they have done: number of inspections performed; number of notices issued; number of formal prosecutions taken, conviction rates and levels of fines imposed. What they have been less good at is evaluating the effectiveness of any of this activity. Moreover, focusing on formal enforcement actions alone leaves a significant swathe of regulatory work uncounted. Yet inspections do not have to result in a formal enforcement action in order to be effective. Giving advice and information can be as valuable as issuing a notice, often more so.

Moving away from counting inputs to evaluating outcomes is a task that no regulator feels that it has yet managed to accomplish successfully. Regulators in different sectors to an extent face quite different problems of evaluation. Environmental or health and safety regulators have the advantage of a large database, and an environment which can be measured. It is relatively straightforward to measure pollution levels or discharges into water, or injury and fatality rates, even if it is difficult to establish a causal link between the agency's action and any increase or reduction in those levels. Regulators in the financial sector face a slightly different problem. They often have to measure invisibles: what would have happened had they not intervened, yet it is difficult to assess a counter-factual. All regulators face the difficulty of knowing when to assess, and how to establish the causal relationship between what they find and what they have done.

Regulators are experimenting with different modes of evaluation, nonetheless, and moving towards more outcome orientated evaluations. The UK Food Standards Agency is moving away from performance targets and reporting based on number of inspections and specified forms of intervention to an outcome based policy focusing on compliance rates. The Environment Agency has set targets for improving operators' management systems based on OPRA scores. The HSE, liberated from input and output targets set by its parent Department, has also moved to assessing outcomes measured in terms of reductions in injuries and fatalities.

Other regulators, notably in the financial sector, have introduced attempts to assess their frameworks in a number of other ways. They look at the movement of firms between risk categories; the regulators' response time to market activities, and stress testing. Stress testing and scenario analysis are used to estimate how firms would cope if certain events were to occur. Six months or so later the regulator will look at whether any of those events did happen and will then compare it with what it thought would happen. Such an approach is however only useful if the management of the firm is relatively stable. In the food industry, where management changes are frequent, such techniques are not as helpful, as the management in charge of the firm can have changed completely since the initial scenario analysis was performed.

Evaluation is important, and the methods by which the regulator is itself evaluated can be in tension with the operation of a risk-based framework. Essentially, what is counted is what gets done. If legislators impose tight restrictions on what it is they will count in evaluating the agency, then they can unduly hinder the regulator's activity and potential effectiveness. In the food sector, for example, EC legislation stipulating the types of "controls" that competent authorities must impose to ensure food safety has recognised only inspections, sampling and analysis and verification of written documents (EC Directive 89/397). The definition of "controls" was expanded in 2004 to include "any other activity required to ensure that the

objectives of [the] Regulation are met (EC Regulation 882/2204 Article 10), and indeed requires controls to be risk-based. This expansion in the types of controls permitted under EC law, and thus recognised by the Commission as constituting a valid control in evaluating member states" food safety regimes, has enabled the UK Food Standards Authority to broaden the types of intervention that it will include in its assessment of local authority food regulation. This has in turn facilitated the development of a new, broader focus on strategies for improving compliance, and indeed enabled the Food Standards Agency's own shift to an outcome based mode of evaluation of local authorities' enforcement activities by requiring them to assess improvements in compliance. There are other examples of where changes in modes of evaluation facilitate the development of outcomes-based policies. As noted above, changes in the evaluation targets for the HSE from inputs and outputs to outcomes has enabled it to move to evaluating its own work in terms of outcomes.

Main challenges of risk-based frameworks

Risk-based frameworks pose particular challenges. The research identifies nine challenges which are of key relevance for those seeking to introduce risk-based frameworks.

Combining simplicity with complexity. Many regulators spoke of the challenge of designing a system which is sufficiently complex to be able to capture and assess a wide range of risks at the firm specific and generic level and which can operate across a widely varying regulated population, and yet be simple enough to be understood used on a day to day basis by inspectors and supervisors.

Knowledge and data. Getting the right data, and making better use of the knowledge the agency has is a critical challenge. Data issues arise both with respect to individual firms and the identification and integration of system wide risks and risks in the external environment which can impact on firms.

Ensuring that assessments of firms are forward looking. Risk assessments often only capture the risks apparent today. Some regulators, such as OSFI, include a "direction of travel" indicator in their risk assessment: is the firm likely to improve or deteriorate over the period to the next inspection? However many other regulators do not explicitly require this assessment, and have found that supervisors or inspectors tend to focus on the risks as they appear now, and not on what might happen in the near future. As noted above, it can also be challenging to ensure that supervisors understand the difference between risk and control – that what they see as a control may in fact be a risk.

Going beyond the individual firm in assessing risk. Here there are two challenges. First, it can be difficult to ensure the framework integrates "horizon" scanning and generic, industry wide risk assessments into the firm-specific assessment. Second, where the regulator has a broad remit, it can be challenging to develop a portfolio approach which compares risks across the whole of the regulator's portfolio of regulated firms, rather than one which focuses on individual firms alone. A single data base on which all firms are scored commonly is critical to effective management across a diverse portfolio, but not necessarily sufficient.

Structure and operation of internal risk governance processes. How to balance the need for organisational structures to ensure the accuracy and consistency of assessments with speed and responsiveness. It is challenging to achieve the right balance between having sufficient internal controls and review to ensure consistency and a hugely bureaucratic framework that in effect stymies the process. It is also difficult to find the right balance

between central direction and local flexibility: allowing sufficient flexibility for supervisors and inspectors to exercise their own judgement, whilst ensuring an acceptable level of quality and consistency of judgments.

Changing the culture to embed the risk-based approach across the whole organisation, from the Board down to individual supervisors. Experience of those who have had risk-based systems for many years suggests that it can take over two years for inspectors or supervisors to really change their approaches and come round. It can take the same time or longer for senior management to really understand the implications of the approach. In some organisations, senior management treated the introduction of the risk-based framework as something that the organisation had to have, but which was not central to what the organisation was doing. As a result there can be a disconnect between what the senior management were doing and what staff were doing.

Ensuring internal compliance with the risk-based regime. Culture changes take time, and a regulator can have a good risk-based framework in theory, but it can be poorly implemented. Developing internal assurance systems to ensure that supervisors and their senior managers are implementing the framework can therefore be necessary.

Managing blame. Risk-based regulation requires the organisation to take risks. A key part of changing culture can be the need to manage blame within the organisation when things go wrong, otherwise supervisors will never feel that they can leave apparently "low risk" issues alone. In non-zero failure regimes, it can be a challenge to resolve the tension between an *ex ante* non-zero failure policy and *ex post* tendency to blame for failures. As one regulator commented, non-zero failure all very well as long as it's not your failure. Senior management support and understanding of the implications of adopting a risk-based approach is thus essential.

Making resources follow risks. There are four issues here. First, resources cannot always track risks with any granularity. Whilst risk-based frameworks can help identify "blocks" of firms, regulators find it difficult to know how to manage resources within the "high risk" block, particularly when they do not have the resources to adopt an intense supervisory relationship or high frequency of inspections with respect to all high risk firms. Second, it is difficult to determine what the appropriate level of baseline intervention should be for the low risk firms. Third, in many regulators the inspection cycle is planned a year in advance; there is always then a lag between risk identification and response. Fourth, the emergence of a politically salient issue immediately diverts resources to dealing with that issue, even if it would otherwise count as low risk, and therefore low priority.

Managing political risk. Politics is often a key driver of what the regulator does. Some regulators seek to manage political risk by incorporating it into their frameworks, or by allowing local flexibility in the allocation of resources to accommodate local concerns. Others, such as VROM, manage political risk by negotiating closely with the relevant Minister as to what their priorities will be for the coming year, and gaining explicit Ministerial approval for their approach. For most, however, their carefully crafted risk-based frameworks are abandoned when politics intervenes. Quite simply, there are some issues with respect the regulator, or the political system, cannot be seen to fail; and it is to those issues where resources will ultimately go.

Conclusion

One of the purposes of the research was to identify lessons which can be learnt from those who have embarked on using risk-based frameworks. The main lessons coming from the research that are of relevance to others are the following (see also IOPS, 2007):

Start with risks not rules. The legislative provisions which a regulator has to implement are often complex and over-whelming. A risk-based approach requires regulators to focus on the risks they need to manage, not the rules they have to ensure compliance with.

Ensure the organisation has sufficient powers to implement the approach. Many regulators have been hampered by inadequate legislative regimes. Regulators need powers to collect the relevant data, and to adopt a flexible approach to determining their inspection policies, and to have a sufficiently wide range of intervention powers. Overly prescriptive evaluation and audit regimes can have similar restrictive effects.

Beware of other regulatory or governmental policies which may contradict or hinder the adoption of a risk-based approach. The impacts of different types of evaluation were noted above. Charging regimes which require regulators to recover the cost of inspections can also distort a risk-based approach to inspection planning. A further example of the unintended consequences are the requirements that those tendering for public sector contracts give details of all enforcement actions. This has been one of the factors prompting companies to dispute enforcement actions taken by the HSE, increasing the resources that it has to spend on investigating and prosecuting accidents as opposed to performing inspections.

Designing and implementing a risk-based framework will take time. As many commented "don't expect it to be right first time". As another observed, "just because something goes wrong doesn't mean the whole system is wrong". Risk-based frameworks are often "built in the lab" by specialists and consultants, and need refinement and adjustment when put into practice. Regulators who are embarking on forming risk-based inspections systems can by now benefit from the experience of others. Nevertheless, pilot projects are recommended by the more experienced regulators in order to trial the framework and to gain "buy in" from firms. If developing from scratch, those how have been through the process recommend that frameworks are developed alongside the on-site inspection process to make sure the two systems match up.

Keep it simple to use and be prepared for the need to make continual adjustments. Frameworks have to be dynamic. They therefore have to be flexible and regulators have to continually revise and update the risk-based model in order to prevent it stagnating and becoming out of step. Frameworks have to be simple to ensure that they are understood by inspectors, and therefore appropriately used by them. In evaluating the framework, use feedback from as many different sources as possible: firms, supervisors, other stakeholders along with internal evaluations to enhance and refine the framework.

Don't underestimate the organisational challenges involved. Organisational challenges are significant, both in terms of changes needed to internal organisational structures and to the changes in skills and culture that will be needed; this may require turnover in staff and a hiring of staff with a different skill set from that sought by the organisation in the past. Systems that were not easily accepted were those that were associated with failures or as having come out of failures, or as being associated with one regulator in particular from a previous regulatory regime (where a number had merged to form the new regulator). In contrast, frameworks were more readily accepted in organisations where they marked

a step change in approach as part of an organisation-wide recognition for a need for change. Nonetheless, it can still take a considerable amount of time for supervisors and senior managers to understand the implications and limitations of a risk-based approach.

Think beyond the risk assessment to how the organisation will respond. There need to be people in the risk assessment process who know what to do when something arises, when the risk crystallises. This can require an integration of people with enforcement experience on the inspection and supervisory teams, and/or a closely integrated compliance and enforcement policy.

Think in terms of achievability. Recognise that resources are likely to be inadequate to adopt an intensive inspection policy even for high risks so think in terms of where those inspection resources are likely to make the biggest difference, and explore alternative strategies to inspections for influencing behaviour.

Communication is vital both within the organisation, with politicians, with firms, and with the public as to what the process is, what the risk scores mean, and how the framework may need to be adjusted. In particular, openness with the industry as to the fact that it is being rated, what the rating means, and that the rating they get will have an influence on how the regulator interacts with them is vital.

It is worth doing. It provides an explicit framework for organising the regulators' assessments and responses. As one regulator commented, "[e]veryone is risk-based and it is better to face up to it and discuss it rather than allowing the organisation to muddle on". Risk-based frameworks can produce resource savings, help to set outcomes and provide a framework for analysing problems or new developments. They can also be used to help set objectives within firms by providing them with assessments of how the firm performs relative to others in the sector, and can provide a common language for discussion with firms' senior management.

But don't do it for the wrong reasons. Learn from others, but don't just adopt someone else's model because people say it is the best. As one regulator commented, "make sure it can work for you; don't adopt it hoping it will miraculously produce a huge internal change, as it won't – that change is very hard to achieve". Recognise its limitations; it is only a tool, and at some point regulators have to ask if it is giving a common sense answer. Risk assessments are inherently judgemental and cannot be purely objective and quantitative, even though many expect them to be.

Recognise that risk-based processes require regulators, and politicians, to take risks, and it is never possible to get consensus on when failures are acceptable. As Douglas and Wildavsky so famously observed, we do not know the risks we face, but we must act as if we do (Douglas and Wildavsky, 1982). Regulators do not know where the next big failure will come from, but they must act as if they do. In so doing, they have to decide whether to err on the side of doing something now that does not need to be done, because it turns out there is no risk; or of not doing something now which it turns out later on that should have been done. Risk-based frameworks can provide a framework for the systematic assessment of political choices, but they can never remove them.

Bibliography

Australian Prudential Regulation Authority (2008), *Probability and Impact Rating System*, APRA, Australia.

Black, J. (2005a), "The Emergence of Risk-based Regulation and the New Public Risk Management in the UK", *Public Law*, pp. 512-549.

Black, J. (2005b), "The Development of Risk-based Regulation in Financial Services: Just 'Modelling Through'?", in J. Black, M. Lodge and M. Thatcher (eds.), *Regulatory Innovation: A Comparative Analysis*, Edward Elgar, Cheltenham.

Black, J. (2006), "Managing Regulatory Risks and Defining the Parameters of Blame: The Case of the Australian Prudential Regulation Authority", *Law and Policy*, pp. 1-27.

Brunner, G., R. Hinz and R. Rocha (eds.) (2008), *Risk-based Supervision of Pension Funds: Emerging Practices and Challenges*, World Bank, Washington DC.

De Nederlandsche Banke (2006), "Focus on Financial Institutions", *DNB Quarterly Bulletin*, June, pp. 71-76.

De Nederlandsche Banke (2006), "DNB's Vision of Supervision in 2006-2010", *DNB Quarterly Bulletin*, December, pp. 54-61.

Department for Business, Enterprise and Regulatory Reform (2007).

Douglas, M. and A. Wildavsky (1982), *Risk and Culture*, UCLA Press, California.

EC 882/2004, "Official Controls performed to ensure the verification of compliance with food and feed law, animal health and animal welfare rules".

Environment Agency (2000), *Delivering for the Environment: A 21st Century Approach to Regulation*, Environment Agency, London.

Environment Agency (2003), *Operation and Management Assessment, Emissions to Air, Guidance on Undertaking an OMA Assessment*, Environment Agency, London.

Environment Agency (2004), *Corporate Plan 2004-2007*, Environment Agency, London.

Environment Agency (2007), *Consultation Paper: Risk-based Regulation of Discharges to Water, Encouraging Better Performance by Businesses*, Environment Agency, London.

Environment Agency (2008), *Greener Businesses, Healthier Environment; Consultation on Regulating Discharges to Water, Response Document*, Environment Agency, London.

Environmental and Spatial Planning Inspectorate (2008), *Evaluation Aspects of the Risk Analysis* (IGAOT), Environmental and Spatial Planning Inspectorate, Lisbon, Portugal.

Environmental Protection Agency (2007), *Code of Practice, Environmental Risk Assessment for Unregulated Waste Disposal Sites*, Environmental Protection Agency, Ireland.

Environmental Protection Agency (2008), *Guidance on Completion of Methodology for Determining Enforcement Categories of Licensees*, Environmental Protection Agency, Ireland.

Fairman, R. and C. Yapp (2005), "Enforced Self Regulation, Prescription and Conceptions of Compliance within Small Businesses: the Impact of Enforcement", Law and Policy, 27(4), p. 491.

Financial Services Authority (2003), *Reasonable Expectations: Regulation in a Non Zero-Failure World*, FSA, London.

Financial Services Authority (2006), *The FSA's Risk Assessment Framework*, FSA, London.

Financial Services Authority (2008), Internal Audit Report on Northern Rock, FSA, London.

Food and Consumer Product Safety Authority (VWA) (2007), *Multi Annual Plan 2001-2011*, VWA, the Netherlands.

Food Standards Agency (2008), *Food Law Code of Practice (England)*, Food Safety Agency, London.

Food Standards Agency (2008b), *Scores on the Doors*.

Food Safety Authority of Ireland, *Code of Practice No. 1 for the Health Service Executive on the Categorisation of Food Businesses*, Food Safety Authority of Ireland, Dublin.

Food Safety Authority of Ireland (2006), *Code of Practice No. 1 for the Health Service Executive on the Risk Categorisation of Food Businesses* (Revision 1), Food Safety Authority of Ireland, Ireland.

FSIS (2007), *The Evolution of Risk-based Inspection*, FSIS, US.

Hampton, P. (2005), *Reduction in Administrative Burdens: Effective Inspection and Enforcement*, HM Treasury, London.

Health and Safety Executive (2004), *A Strategy for Workplace Safety in Great Britain to 2010 and beyond*, Health and Safety Executive, London.

Health and Safety Executive (2005), *Enforcement Management Model*, Operational Version 3.0, Health and Safety Executive, London.

Health and Safety Executive (2008), "Field Operations Directorate Topic Packs", available at *www.hse.gov.uk/foi/internalops/fod/inspect/*.

House of Commons Select Committee on Work and Pensions (2007), *Role of Health and Safety Commission and the Health and Safety Executive in Regulating Workplace Safety*, 3rd Report, Session 2007-8, HC 246, HMSO, London.

Hutter, B. (2005), *The Attractions of Risk-based Regulation: Accounting for the Emergence of Risk Ideas in Regulation*, CARR Discussion Paper No. 33, London, LSE, 2003.

Hutter, B. and S. Lloyd Bostock (2008) "Reforming Regulation of the Medical Profession: The Risks of Risk-based Approaches", *Health, Risk and Society*, 10(1), pp. 69-83.

IMPEL (2007), *Doing the Right Thing II: Step by Step Guide for Planning Environmental Inspections*, IMPEL (European Network for the Implementation and Enforcement of Environmental Law), Brussels.

International Organisation of Pension Fund Supervisors (IOPS) (2007), "Experiences and Challenges with the Introduction of Risk-based Supervision for Pension Funds", *Working Paper*, No. 4, IOPS.

Meyer, J. and B. Rowan (1977), "Institutionalised Organisations: Formal Structure as Myth and Ceremony", *American Journal of Sociology*, 83(2), p. 340.

National Audit Office (2007), *The Pensions Regulator's Progress in Establishing its new Regulatory Approach*, National Audit Office, London.

National Audit Office (2008f), *Regulatory Quality: How Regulators are Implementing the Hampton Vision*, National Audit Office, London.

National Audit Office and Better Regulation Executive (2008a), *Effective Inspection and Enforcement: Implementing the Hampton Vision in the Environment Agency*, NAO and BRE, London.

National Audit Office and Better Regulation Executive (2008b), *Effective Inspection and Enforcement: Implementing the Hampton Vision in the Financial Services Authority*, NAO and BRE, London.

National Audit Office and Better Regulation Executive (2008c), *Effective Inspection and Enforcement: Implementing the Hampton Vision in the Food Standards Agency*, NAO and BRE, London.

National Audit Office and Better Regulation Executive (2008d), *Effective Inspection and Enforcement: Implementing the Hampton Vision in the Health and Safety Executive*, NAO and BRE, London.

National Audit Office and Better Regulation Executive (2008e), *Effective Inspection and Enforcement: Implementing the Hampton Vision in the Office of Fair Trading*, NAO and BRE, London.

Office of Environmental Enforcement (2007), *Guidance on Completion of Methodology for Determining Enforcement Category of Licences*, Environmental Protection Agency, Ireland.

Rothstein, H., M. Huber and G. Gaskell (2006), "A Theory of Risk Colonization: The Spiralling Logics of Societal and Institutional Risk", *Economy and Society*, 35(1), p. 91.

Schrader-Frechette, K.S. (1991), *Risk and Rationality*, University of California Press, Berkeley.

Slorach, S.A. (2008), *Food Safety Risk Management in New Zealand: A Review of the New Zealand Food Safety Authority's Risk Management Framework and its Application*, NZFSA, New Zealand.

Sparrow, K. (2002), *The Regulatory Craft*, Brookings Institute, Washington DC.

The Pensions Regulator (2006), *Medium Term Strategy*, The Pensions Regulator, London.

The Pensions Regulator (2007), *Corporate Plan 2008-2011*, The Pensions Regulator, London.

VROM (2004), *Compliance Strategy for the Ministry of Housing, Spatial Planning and the Environment*, VROM, the Netherlands.