CHAPTER ONE

Introduction to Empowered Enterprise Risk Management

A CLEAR SIGN THAT YOU are a helicopter parent, according to online sources, is developing a bad back from constantly stooping down and following your toddler's every step. Helicopter parents are, as it were, those constantly trying to identify and remove threats to their child's safety. They hover above the playground, ready to interfere at a moment's notice, and generally put a variety of restrictions on the child's activities to remove any notion of danger. Such a highly regimented style of parenting is in sharp contrast with the much more relaxed approach that was common not too long ago. As recently as the 1980s, even in that epitome of the safety-first approach, in Sweden it was not uncommon to see small children standing up between the front seats of the car while the car was being driven. Their parents would not necessarily have been viewed as irresponsible by other adults or reflected much themselves on the possibility that they might have been taking unacceptable risks.

The helicopter parent is just one caricature describing a broader current in society, namely a desire towards identifying and controlling risks that might affect our well-being. Sociologists have even referred to our modern world as a 'risk society', meaning a society that is increasingly preoccupied with the future and any risks that it might bring (Beck, 1992). It turns out that modernity has ushered in a wide variety of man-made risks on top of the natural

hazards that have always threatened societies. The complexity of the systems that support modern life, and the degree of their interconnectivity, have generated a whole range of new risks. At the same time, there has been a veritable explosion in the availability of information, making us ever more aware about potential threats. New norms emerging over time have gradually come to present an attitude of caution as a necessity.

The growing focus on risks has been accompanied by a belief that not only should they be managed but *can* be. As Peter Bernstein's epic story of risk shows, over the millennia we have gone through a series of intellectual revolutions, from being 'passive before the gods' – that is, largely resigned to fate – to claiming mastery over risk (Bernstein, 1996). The premise seems to be that science and our expanding knowledge can be used to assist us in more safely navigating the world. This premise has given rise to an entirely new profession: risk management. In its early applications, the position of risk manager typically referred to a highly specific area of expertise. However, ideas about risk management have reached further and further up the corporate hierarchy, infiltrating even top management teams and boards of directors. Power (2007) writes:

> In a relatively short period of time, in a number of different countries, hospitals, schools, universities, and many other public organizations, including the very highest levels of central government, have all been transformed to varying degrees by discourses about risk and its possible management.

Risk management, in this view, has gone from being a specialized subfield to being a source of principles for organizing and managing in general. To put it another way, there has been a shift from risk analysis, a technical discipline, to the *governance* of risk in organizations.

Interest in risk governance took a giant leap in the early 1990s with the publication of two reports that were to become highly influential: The Cadbury Code, published in the UK in 1992, and the COSO framework for internal control, also published in 1992.[1] These reports contain a set of recommendations for achieving sound governance in organizations, and to generally improve oversight over vital processes. The background for both reports was a string

[1]COSO (Committee of Sponsoring Organizations of the Treadway Commission) is an 'umbrella organization' consisting of the following five organizations: Institute of Internal Auditors; the Association of Accountants and Finance Professionals in Business; Financial Executives International; American Institute of CPAs; and American Accounting Association.

of corporate failures that were deemed to have been triggered not so much by exogenous risks or flawed analysis, but rather failures of governance. In other words, there was a systematic lack of checks and balances in organizations that left them too vulnerable to fraud and other misbehaviours. To improve this situation, boards of directors were encouraged to, among other things, put in place systems designed to increase control. This led to the codification of what is known as 'internal control', defined in COSO (1992) as follows:

> Internal control is a process effected by an entity's board of directors, management and other personnel designed to provide reasonable assurance of the achievement of objectives in the following categories:
>
> Operational Effectiveness and Efficiency
> Financial Reporting Reliability
> Applicable Laws and Regulations Compliance

Internal control can be viewed as an effort to address the control problems that afflict organizations without adding new and costly external regulations (Power, 2007). Instead, risk control was to be achieved from within, through 'self-discovery and reporting'. To this end, new functions like internal audit and compliance were set up in many firms, in no small measures boosted by the arrival of the Sarbanes–Oxley Act in 2002. Taking stock some thirty years later, the focus on internal control can certainly be said to have had a lasting impact on business practices. True, the application of the concept has been varied and frequently modified to reflect firms' specific needs and capabilities, but the core idea of maintaining an independent function to safeguard the integrity of important processes has stood the test of time.

The introduction of new functions related to internal control was a major leap in the broader trend towards organized forms of risk management. One of the five pillars of COSO's framework is that internal control should be 'risk based', which is to say, it should be preceded by an inventory and assessment of the risks that could pose a threat to the entity's objectives. However, the reach and impact of risk management ideas has continued unabated. Contributing to this development was the fact that internal control was found to be lacking in two main respects. One was that it was not perceived as holistic, or enterprise wide, enough. The risks that fall within its scope are only a subset of all the different kinds of risk that businesses face. There are market risks, reputation risk, business disruption risk, and so on, in an almost endless variety. The second factor that spurred on further development was the feeling that

risk management ought to provide more active support to business decisions, thereby working closer with management and 'adding value'.

Functions dedicated to internal control are held back in such pursuits by their pledge to remain independent vis-à-vis management. Internal control, at its core, chiefly seeks to contain threats. A different set of ideas would be needed to realize the vision of an enterprise-wide form of risk management that supports business strategies. The principles behind holistic risk management were first developed in the mid-1990s by the people behind the Australian risk management standard (AS/NZ 4360) with later additions by their Canadian counterparts. But again, it was COSO that delivered the blueprint that was to shape much of the field's future development. In 2004 the Enterprise Risk Management (ERM): Integrated Framework was published. It defines ERM as follows:

> Enterprise risk management is a process, effected by an entity's board of directors, management and other personnel, applied in strategy setting and across the enterprise, designed to identify potential events that may affect the entity, and manage risk to be within its risk appetite, to provide reasonable assurance regarding the achievement of entity objectives.

In this definition the ambitious agenda of ERM is plain to see. It is no longer just about risk control, but it shall support decision-making in a 'strategic setting'. Risk management is now to be carried out across the enterprise and is not limited to compliance, financial reporting, and operating hazards. Another game changer was the introduction of the term 'risk appetite', which signalled that some level of risk can be tolerated. In fact, the COSO (2004) text makes explicit mention of the concept of 'capitalising on opportunities', a perspective that is absent in the world of internal control. In 2017, COSO published, jointly with consulting firm PWC, an updated version of its framework that emphasized even more strongly the connection between strategy and ERM. The other main risk management standard, ISO 31000, issued by the International Organization for Standardization, shows a similar trajectory towards a strategic agenda involving senior management.

For all these advances in the thinking about the role of risk management in firms, it is significant that its most cited blueprint grew out of the world of internal control. Lest nobody mistakes where it came from, COSO (2004) even states that ERM builds on, and completely encompasses, the internal control framework. This has meant that the ethos of internal control, with its emphasis on risk control rather than balancing risk and return in a pursuit

of value, has come to have a disproportionate effect on the implementation of ERM. Power (2009) makes a key point: 'The ERM model is strongly, if not exclusively, influenced by accounting and auditing norms of control, with an emphasis on process description and evidence.' He goes on to comment on the proliferation of detailed processes for risk management based on rules and prescriptions: 'Accounting ideals of internal control are embedded in the design itself, resulting in a style of risk management practice with wide and seductively expansive reach – the risk management of everything.'

This leads to a question that is still not fully resolved. Is ERM supposed to be 'just another control function', or an in-house advisory that works closely with the executive team on matters of strategic importance? Many managers indeed seem to associate ERM more with controlling risk than anything else, and therefore seem content to just produce a risk map, check risk management off the list, and carry on as before. As a result, ERM could amount to what Power (2009) refers to as 'the risk management of nothing': a superficial effort that fails to drill deep into the interconnected nature of risks. Any feeling of safety afforded by it is therefore an illusion, or even worse, misleading, because it largely fails to articulate and comprehend critical risks.

In fact, puzzles and paradoxes abound in the COSO definition of ERM. Another example of an unresolved issue is that of the real meaning of risk appetite. The core questions we may ask are:

How much risk can we tolerate?
How exactly is risk supposed to be traded off against upside potential?

The latter balancing act is the dimension that risk appetite is supposed to bring much-needed attention to when compared a framework geared towards risk control. But as so often with these frameworks, they only establish certain basic guiding principles and leave much of the interpretation open to the organizations doing the implementation. Truth be told, however, there are probably few concepts in business that have generated similar levels of confusion as risk appetite. Risk appetite as a concept seems to be almost perfectly designed to do so by appealing to the subjective nature of risk rather than to an analytical, fact-based approach, and for being something of a contradiction in terms. Yet it pervades the discourse on risk management today. And yes, we do need a way to impose a limit on risk-taking and put that in the context of business opportunities. So, how is this dilemma to be resolved?

Another paradox in ERM concerns the role of objectives in determining risk management strategies. In the main frameworks, COSO and ISO 31000,

the achievement of objectives is put on a high pedestal. In practice, many firms have taken this to mean protecting short-term targets, or the so-called key performance indicators (KPIs) that are widely used by firms to measure progress. While protecting targets sounds fair enough, this leaves ample room for behaviours that are inconsistent with the higher-level objective of creating value. In fact, it is not hard to illustrate situations where spending resources to increase the probability of target achievement is detrimental to long-term firm value. And to completely flip the perspective, there is growing evidence that the targets themselves may be a *source* of risk. The performance pressure induced by such targets has been known to cause reckless behaviour and gambling on a scale that can lead to firms getting into serious trouble. The Deepwater Horizon disaster, under the watch of BP, and the fall from grace of US bank Wells Fargo are only two examples of target-chasing as generators of risk.

And what happened to firm value as a higher-level objective anyway? Here is yet another paradox. Most managers find it self-evident that they are in business to generate profits, and ultimately dividends for their shareholders. Treatises on ERM often do speak about shareholder value, but on the whole the emphasis is not that strong. In fact, the ERM frameworks seem to be embedded, to a fair extent, in the *stakeholder* value paradigm. The first line in COSO's executive summary (2004) reads as follows: 'The underlying premise of enterprise risk management is that every entity exists to provide value for its *stakeholders*' (emphasis added). In stakeholder theory, firms exist to satisfy the interests of their many stakeholders, such as employees, customers, the state, and so on. Shareholders are just one of many stakeholders and have no special privilege when it comes to setting the firm's objectives. When business units, the very risk management 'silos' that ERM was meant to integrate, discuss their risk appetite, from a stakeholder perspective, with respect to the risk of failing to meet short-run targets, we have a pretty confused situation.

The paradoxes do not end here. We have also observed, among ERM practitioners, a curious reluctance to quantify and perform financial analysis. Another observation is that often the responsibility for a risk is simply delegated to the person who is already working in that part of the business. With the result that some sort of governance of risk has taken place, at least on paper, but nothing of substance seems to be elevated to a *centralized* form of management, as defined in the portfolio view of risk underpinning ERM. What is more, for all the talk about holistic, enterprise-wide risk management, ERM programmes seem generally incapable of addressing questions related to total risk. Any notion of what the risk of the firm is as a whole continues to elude firms that implement ERM, to the extent that this perspective is not even brought up at all.

Amidst all these paradoxes and puzzles of ERM, there is the nagging concern that, for a framework that loves to emphasize the need to use risk management as a tool in the pursuit of upside and business opportunities, there seems to be a lot of 'covering your back' going on. Is it possible that the result of all the governance taking place, and the processes needed to support it, is that firms are more anxious and slower to respond to new opportunities? Formal risk appetite statements written into policy documents, for example, do have a whiff of corporate bureaucracy attached to them. Just like helicopter parents may be accused of stifling their children's autonomy and independent judgement, could the apparatus around ERM contribute to the dulling of the entrepreneurial spirit in firms? Do people become so worried about making errors that they consistently prefer to stay within their comfort zones instead of venturing out into new and risky, but also potentially very lucrative, projects? This is certainly not the intended meaning of either the internal control or risk management frameworks. But it suggests that how such organized forms of risk management present themselves – their 'vibe' – is quite important. Is it going to be another round of control and compliance, or is it a partner in business, an enabler that that helps the firm compete, succeed, and create value? In this book, we support the latter version of ERM. While more difficult to achieve, the benefits are also far greater.

WHY A THEORETICAL PERSPECTIVE?

ERM has obviously moved on since the COSO (2004) document was published. ERM is sometimes described as an evolving phenomenon, which may take many years before it becomes codified and practised in a consistent way. In this view there is an ongoing search for best practices that eventually will settle into a body of concepts and practices that will constitute ERM. The idea is that best practices will evolve, in a Darwinian-like manner, if given enough time. Along these lines, Mikes and Kaplan (2015) argue that there is no one universal form of ERM that will be right for all firms. Rather, each firm chooses from the available design parameters to obtain an 'ERM-mix' that is suitable to its particular circumstances.

While there is certainly something to be said for letting robust practices evolve by proving their usefulness in actual practice, we also take a somewhat different view. We believe that ERM stands to benefit from a more rigorous description, at the theoretical level, of the problems it is supposed to solve in the first place. The definitions of ERM provided by the frameworks are just that – definitions. And definitions are not theory, except by accident.

They are more like opinions, descriptions of how one sees the world or would like things to be. Reflecting the endless malleability of ERM, there are indeed definitions aplenty on offer, not just in the frameworks but from large numbers of authors who have their own take on ERM. Theory, in contrast, draws attention to the root causes of the problems afflicting practice, and can therefore point to the appropriate solutions. It provides focus and structure to the design of ERM. We are not just 'doing risk management', but rather addressing certain well-defined problems with the goal of improving decision-making in pursuit of a defined higher-level objective. ERM, in this more theoretical approach, is derived from a set of first principles instead of just conjured up from definitions.

A more theoretical view of risk management in firms starts with some basic premises about what is being optimized – the ultimate goal. This is usually taken to mean firm value. In the classic theory of corporate finance by Modigliani and Miller (1958), however, capital structure (and risk management by implication) turns out to be irrelevant. The reason underlying this somewhat disturbing result is that the authors make several very strong assumptions – such as no taxes, a fixed investment opportunity set, and equally distributed information. Much of the theory since this seminal work has progressed by investigating what happens to optimal corporate policy if one or more of these assumptions *fail* to hold. In fact, almost none of the assumptions bear scrutiny in the real world. Obviously, it is fairly indisputable that taxes and bankruptcy costs exist in the real world. But we also know that decision makers suffer from behavioural biases, conflicts of interest, and lack of complete information.

The optimal corporate policy is, broadly speaking, the one that minimizes the impact of all these imperfections on firm value. The use of financial derivatives to manage risk ('hedging') and insurance are two of the policies that have been investigated from the perspective of minimizing the impact of these frictions. These strands of research amount to what could be referred to as 'classic' risk management theory, which has delivered many important insights that we will come back to numerous times in this book. Hedging to reduce expected costs of bankruptcy, for example, is a long-standing example of how the use of derivatives can increase firm value.

In these academic models, the firm itself is characterized as a *unified entity* interacting with providers of financing who know less about the company's prospects (information asymmetries) and suspect that managers have hidden agendas (conflicts of interest). Just like in classic economic theory in general, it is usually taken for granted that whoever sets policy in the firm has access to full information, and that it is the ones on the outside that struggle with lack of

information. If there is a conflict of interest, it is between the firm's managers and the investors in the company, not between different layers of the firm.

Reality overwhelmingly suggests, however, that firms are not unified entities. Even if the executive team's interests have been aligned with those of shareholders (which we take to mean maximizing firm value), there remains the issue of how risk management can be applied for this purpose in an enterprise consisting of several business units with decentralized decision-making, for whom firm value is a distant and abstract concept. History is also replete with examples of executives and directors who were not able to understand the firm's exposures to risk and were taken by surprise when something bad happened. In surveys, executives regularly point to challenges in aggregating information about exposures, and in increasing the 'visibility' of certain risks.

In our analysis of ERM we shift the attention to how some of the aforementioned imperfections operate *within the firm*. ERM is much more about the risk management *process* than any specific risk management *strategy*. It belongs in the realm of corporate governance and management control, both needed to govern decision-making processes that are at least partly decentralized. The focus will be on behavioural biases, conflicts of interest, and information asymmetries as they play out on multiple levels in a firm with decentralized decision-making. ERM thus takes place at the interface between the board of directors, the executive team, and multiple business units and corporate functions. Does the board of directors really have access to full information about the main risks? Do business unit managers really use risk management to maximize value or do they pursue other agendas? Are business units too optimistic and willing to pursue ventures that imply too much risk? In contrast to the view that ERM is an evolving set of techniques, to be chosen by each firm from some kind of smorgasbord, these behavioural, agency, and information issues are themes that should guide the design of risk management in any firm.

Theory, then, is useful because it articulates with greater clarity the problems that ERM is meant to solve, pointing us to the root causes of inferior risk management execution. Our interest lies in using theoretical perspectives to understand the problems at hand, because we believe doing so simultaneously points to the most effective solution. In this book the reader will therefore find no complicated math or excessive jargon for its own sake. To us, theory is basically only a set of well-supported ideas about how the world works. We strongly believe in keeping things simple, and that the real point is to improve thought processes (and by extension decision-making). While our exposé covers some fairly advanced topics, our basic message is that most of the benefits can be reaped through a few relatively simple changes in mindsets and practices that

are well within the reach of all organizations, whether small or large and regardless of industry.

LESSONS FROM PRACTICE

However, as we pointed out in the Preface, theory is not enough. We also need to better understand how risk management principles can be embedded in organizations and achieve impact. ERM wants to make a difference in practice, yet often comes up against powerful forces that seek to limit its influence. ERM calls on people, at least to some degree, to change their way of thinking and how they do things, and anyone who has worked in an organization knows that this is no small task. Therefore, ERM initiatives often have an element of change management to them. ERM in our case company, Equinor, has even been described by company insiders as a 'culture project' as much as being about specific tools and techniques.

What we can learn from studying practice is, above all, what the keys to successful implementation are (and the traps to avoid). What are the things that need to be said and done to overcome the obstacles and resistance to change? What resonates with people at different levels in an organization and motivates them to take on ERM? What makes them strive together towards a common goal rather than be stuck in opposition? All in all, we want to find out about what really works and what helps unlock the potential of ERM. Theory cannot inform us a lot about these things. One has to get out in the midst of organizational life and the constant to and fro of different ideas, trends, power struggles, individual quirks, and so on. But for all the variety and richness of this setting, certain patterns are more likely to be associated with success than others. Thus, we hope and believe that many of the keys and success factors uncovered in our empirical material are generally applicable.

What is a good case company for an in-depth exploration of ERM in practice? We believe that the answer to that question is a company where ERM has demonstrably been impactful and where it has been practised for many years, interlocking ever more deeply with decision-making processes. Equinor meets these two criteria. It was one of the early adopters, launching an ERM programme in 1996. ERM has been continuously used and refined since then, from an initial focus on market risks to a fully enterprise-wide effort that actively engages the board of directors as well as smaller operating segments. Equinor has indeed reached the point where risk management is considered a core value, or 'a way of life', as it is sometimes called. It is not merely

a technical exercise by specialists in the risk management function. Understanding and managing risk is something that is practised by – indeed expected from – each employee.

A case study on how Equinor got to this point holds the promise of offering insights into how ERM can be best approached in practice. We have tried to refine into a compact and accessible format all that Equinor has learnt the long and hard way. The ambition is to make the journey for someone just starting out quicker and more direct to the desired destination. What awaits at the end of that journey is empowered ERM, which occurs when the organization commits to risk management and elevates it to a core value, central to its way of doing things and how it defines itself. As a result of this commitment to proactive risk management, there will be better information and conversations around risk; clearer responsibilities; a unified language and methodology; and improved business decision-making that takes due account of risk.