



# Introduction

*You cannot expect the form before the idea, for they will come into being together.*

Arnold Schönberg

## 1.1 Why this book?

Financial services institutions such as internationally operating banks and insurance companies frequently need to adapt to changes in their environments, yet at the same time manage risk and ensure regulatory compliance. This book explains how metadata is key to managing these goals effectively.

In recent years, an array of scandals has adversely affected the trust that investors, regulators, and the general public put into reports from corporations, specifically when it comes to financial accounting. Some of these scandals can be attributed to the exploitation of legal loopholes that eventually backfired, others were caused by plain old fraud and criminal energy to cover balance sheet irregularities. And yet, there were cases where people discovered that they had made, well, an honest mistake. Whatever the reason, an increasing body of legislation puts an increasing responsibility on companies to deliver accurate and consistent reports. The need for ensuring compliance end-to-end is on the up as well, and solid, systematic corporate governance has become a necessity.

In addition, worries about the stability of the global financial system caused national and international regulators to ponder further standardization, to reduce the risks inherent to the system by increasing the comparability and transparency of risk management methods and procedures, as well as to set limits for compliance. Some of these standards are driven by a

desire to increase market transparency, whereas others reinforce requirements to avoid conflicts of interest, intended to reflect the improved understanding of a specific subject, or focus on market, credit, or operational risks. Regulations and frameworks like SOX, COSO, Solvency II, MiFID, IFRS/IAS, ITIL, COBIT, or Basel II come to mind. The topic of risk management is high on the agenda. Here again, compliance with regulatory requirements and corporate governance is a prominent concern.

All these obligations have to be met by businesses operating in the financial services industry, which has been particularly affected by these developments. At the same time, this industry has also witnessed a steady increase in the complexity and sophistication of its products, services, and processes, as well as the markets within which it trades. It does not exactly help that managing such a zoo is subject to a vast lineup of regulations formulated by authorities.

These challenges can be considered to be a regular side effect of maturation and industrialization. The past has also witnessed market participants engage in risks, get their fingers burnt, learn from it, and move on. After all, an important part of financial services is about engaging in and managing risks. However, a particular concern about the current business climate is the sometimes drastic effect of misbehavior or misjudgment. Investor confidence can plunge without warning, so the worry about reputational risk has grown to extremes in some areas. It is therefore no surprise that the topic of governance and risk is a concern.

Another factor to consider is the continuing push of the financial services industry towards integration and globalization of their business. Some areas are already globalized (think of investment banking or reinsurance), and others are catching up. Mergers and acquisitions will be increasingly transnational, crossing jurisdictions and markets. This raises issues that were traditionally less of a concern. A bank operating across jurisdictions that has not done so before must deal with a much wider range of constraints and comply with a larger number of rules than a national niche player. In some regions, life has been made easier by unifying the rules under which companies can operate cross-border, for example the European Union. But that offers little comfort to the businesses that truly want to compete globally.

Complexity, variability, volatility, change, cost pressures, constraints, and risks abound in financial services. The goal conflict is evident: on the one hand, the call is for increased flexibility and agility to support the change that companies are undergoing; on the other hand, wide-ranging control over the risks run is required, and the alignment of organizations with legislation (and, not to forget, stated regulation internally) is demanded by company leaders even more strongly. What is more, besides regulatory requirements there are other sources that may require you to adapt to change, such as competition, trends, or innovation.

It would be unhealthy to expect one single solution to resolve this goal conflict completely, because it is systemic: the more risks you run, the more likely you are to witness volatility in your results. If you allow employees to act flexibly, you cannot realistically expect them to comply with all existing rules, unless you check compliance with them... which costs your

employees parts of the very flexibility you just gave them. The more complex your products are, the less you will be able to understand the risks attached. However, as these conflicts can potentially translate into direct financial gains or losses, you will be concerned with balancing them wisely. Managing this balance well gives you the ability to put your capital to the best possible use.

What comes in handy at this point is the fact that the business processes of financial services companies depend – to a large extent crucially – on information technology (IT). This offers an opportunity to use leverage effects (digitization, automation, scalability) for the benefits of supporting change, which can be understood as a business process itself. Yet, although leverage is generally a fine idea, one should be aware of the fact that too much bureaucracy and risk (into which highly automated processes can easily degenerate) can be as much of a disaster as too little.

Dealing with the above goal conflict (flexibility and speed on the one hand and risk and compliance on the other) is thus about dealing with it in a constructive, systematic fashion, while confining undesired side effects to areas that are of lesser importance or, better still, no concern at all. This is what I want to call *systematic alignment*, which means that:

- Change is supported in a way that increases the effectiveness of moving the corporation from an existing to a desired target state.
- Effectiveness is increased by giving structure to a change and leveraging this structure, using information technology, to arrive at the target state in a controlled fashion.
- The nature of a change and its impact is used for alignment with regulations, which are formulated as constraints to check compliance with the target state. It is also used to understand the risk characteristics of the existing and target state.
- Processes are put in place at each of these stages to plan, execute, and monitor goal achievement, while ensuring the necessary flexibility to cater for the complexity arising in real-life settings.

This book is about the data to support systematic alignment, called *metadata*, and the processes to handle it, called *metadata management*. It explains how the use of metadata can help you add value by improving performance, managing risk, and ensuring compliance. It is specifically a book for the financial services industry, its peculiarities, needs, and rules. It elaborates where and why metadata helps this industry to manage change.

In a word, metadata is a choice for people who want to transform their financial institution in a controlled fashion, and yet be sure they comply with regulations and control their exposure to risk. Managing this conflict in a structured, predictable fashion can be made the centerpiece of managing change successfully, confidently, and reliably. But why, again, should that worry you in the first place?

## 1.2 Change, risk, and compliance

In October 2003, a US Securities and Exchange Commission official cited the following key examination and enforcement issues that the SEC worried about:

1. late trading and timing in mutual fund shares;
2. creating and marketing structured finance products;
3. risk management and internal controls;
4. accuracy and reliability of books, records, and computations.

It is probably safe to assume that much of this continues to be on the agenda. The complexity and sophistication of products and services offered by the financial services industry continues to increase. As a consequence, the change frequency will also increase, and the impact of changes will become more difficult to assess, and less easy to predict. That is, of course, unless countermeasures are taken. Of the above concerns, three out of four can – at least to a substantial degree – be attributed to complexity, rapid change, and intransparency.

The reasons for this are easily explained: products and services tend to differentiate, be packaged and configured to the needs of the customers, and structured to appeal to more complex demands – in a word, evolve. Structuring is one such typical practice: as products become commoditized, they are used for putting together other products to arrive at more powerful, yet subsequently more complex products. On the other hand, the number of dependencies to be managed increases, causing management problems of its own kind.

This combination of scale, dependency, and change is a substantial source of risk. Typically you would want to carve up this cake and opt for some form of divide and conquer. But how, where, and when? There may literally be dozens of factors you need to take into account. Not all of these factors will you be able to influence, not all of them may be worth your full attention, and not all of them lend themselves to carving up, anyway.

Think about this: the typical financial institution operating on a global basis will sport a portfolio of applications numbering in the upper hundreds, not counting different releases of these applications, desktop software, spreadsheets, and the like. You will naturally be disinclined to take on the full scale of the problem just yet. Hence, you decide to put out of scope, spreadsheets and desktop software. But there the problem begins: your risk managers use a combination of spreadsheets and the desktop installation of a risk modeling and simulation package. From the relevance viewpoint this should certainly be in scope, but from the scalability viewpoint you will not want to include all spreadsheets and desktop software packages. To mix and match at will, however, just worsens your situation once again: how are you going to draw any meaningful conclusions about a bag of seemingly incomparable things?

The answer depends, to a large extent, on your perspective. Which perspective you take, which aspects of your problem you *reify* into a model to treat in a structured fashion, has a

great influence on what you will be better at, but also where you will, well, not be better. The question is, which aspects should you focus on, and what are the undesirable effects of that choice? And, should these aspects constitute different, competing views on the parts of the same problem, how can you balance them in a meaningful way?

In a word, if you want to manage risk in the face of scale, change, and dependency, you need to find a way to structure it. What you also need to consider is the rate and impact of change. What does this mean for the IT applications of a large company? What are the conditions under which the business processes of such typical companies operate? What are the origins of change? Can the risk this causes be controlled? Where? And when? How can IT support this process? It is here where the topic of systematic alignment comes up again: many business processes are run by applications (partially or completely). Hence, there is substantial leverage you can exert over the risks associated with the execution of these processes by exploiting the structure intrinsic to any software product.

Take another example, financial accounting. Traditionally a complex mix of formality and fluidity – of rigid rules and flexible interpretation – this area is complexity galore: the general ledger of any larger institution can easily contain hundreds of accounts, each of which features its own list of describing attributes. But there are also the dozens of subsidiary ledgers, based upon which the general ledger is computed, and there are its numerous variations, all of which are governed by their own set of rules (both written and unwritten) and levels of detail: report to analysts, quarterly report, annual report, and so on. Chip in different accounting standards in different countries, and then imagine migrating from US-GAAP to IFRS/IAS-based accounting. Furthermore, assume your company is listed on a US stock exchange, subjecting it to the regulations of SOX. You will be bound by those rules as well, most notably those of the (almost notorious) Section 404.

Most certainly you want to make sure all these rules are abided by, but would you know how to ensure they really are? The internal control procedures in a large organization are typically complex, often interwoven, changing, and most likely not understood by any one single person in the company at every level of detail. Therefore, you will want to get a grip on understanding, for example, which processes enact changes on accounts considered significant enough to merit control, and then subject those processes to controls. Yet, there are many different ways in which a process can be enacted, e.g. by IT only, by humans supported by IT, or by humans only. As new applications are developed, these ways change, and as new entities are acquired or merged with, the whole fabric is put under tension. Sometimes it tears and needs to be fixed.

The complexity cannot be done away with. Hence, as a measure of caution you will want to give the problem some structure so to understand it better – a framework. Today, many organizations interested in achieving SOX compliance use the COSO framework. In a stable environment, you are certainly safe with this. But the framework does not (directly) answer

the question how to approach the issue of integrating a – possibly large – acquired company, or reorganizing the entire corporation, such that the proverbial fabric does not tear.

This is where systematic alignment comes in: frameworks such as COSO, COBIT, and others take specific structural decisions in order to achieve predefined goals, which can be technical, legal, financial, or something else. They do so by classifying things – be they tangible or not – in a certain way and then establishing rules on their (structural) relationships.

For example, concepts in the world of the Sarbanes-Oxley Act include accounts, risks, and processes; the relationships are ‘transacts on,’ ‘is child of,’ and ‘can affect;’ the rules are ‘aggregate’ and ‘cascade.’ If you want to integrate the other company, its accounts, its processes, its risks, the whole lot must be aligned with the ones you are managing in your own company. Here is where the framework, its concepts, and its structural rules, along with metadata come to your rescue: by mapping the other company’s processes to your own (capturing metadata), you can systematically establish a desired target state (alignment with your own processes) and track it (monitor compliance). Usage of the framework ensures that the business meaning of concepts on either side of the integration process is the same, thus ensuring that the integration does what is intended.

One last example: the creation, marketing, and trading in financial instruments. The markets they are traded in are extremely innovative in coming up with products tailored to very specific needs. These are typically highly sophisticated, and their inner workings are accessible to but a few people who understand the market risk associated with trading in them. It will be genuinely difficult to come up with a risk model for these products. There are situations where you will be incapable of fully understanding a problem, let alone the structure and reason about it. This is despite the fact that financial instruments are often composed of others, thus revealing their structure to the observer. As a matter of fact, in many such situations structuring the risk is not the first choice, but using completely different techniques such as exposure limitation or hedging.

Speed and flexibility are crucial in these markets. The proverbial first-mover advantage can mean money lost or won. As such, it would be flat-out silly to subject the designers of complex financial instruments to a rigid regime of rules, constraints, and regulations on the grounds of principle. Not only would you very likely miss out on some nice business opportunities, but you would also miss out on learning how to get better.

The world of managing risk and compliance for financial instruments is a good example of an *ill-structured* problem, meaning that it is full of complexity, ambiguity, and there is sometimes no obvious specification to describe it formally. Such problems can typically be seen from many different viewpoints and require many dimensions along which to describe them properly. To master them is less a question of structuring savvy but of experience or if you will, expert judgement. This makes them hardly accessible to structuring and formalization. Since the financial markets are moving so quickly it is imperative to avoid smothering your staff in

red tape. You will want to grant them the flexibility they need to be creative, but only to an extent. You are up for a goal conflict.

Actually, there is another side to the story of financial instruments, and it is about commoditization. You may have noted that it was said that many of today's financial instruments do in fact have structure, and that they are often built on top of other existing, widely-available, instruments. So, there must be some modeling and structuring going on. Can we not make use of that? How does that all fit together? The answer lies in the nature of what (and the place where) something is changing and how to support that process.

Right after a financial instrument has been invented, the risk associated with it is typically not yet understood well. The market sways and swings, and market participants learn how to treat it properly. For example, many accounting issues surround insurance-linked securities, which have jurisdictional dependencies (e.g., US-GAAP versus IFRS/IAS accounting). Ultimately, the risk and compliance issues associated with the instrument are thoroughly understood and they become available as an off-the-shelf product. They may even become traded in a financial exchange, at the time of which they have finally become commoditized. As soon as that has happened, new financial instruments can be built on top of it, and the cycle begins anew.

Although the instruments are not yet thoroughly understood, managing risk (or compliance, for that matter) is a try-and-fix exercise. As they mature, standard products emerge, which can be reused in other settings. However, now that risk and compliance is, for all practical purposes, fully understood, their properties can be used constructively. For example, modeling and simulation tools can be built to compose other products on top of them. These tools in turn make use of said (structural) properties, which are metadata. Modeling, simulation, and design, are in turn highly creative activities, and we go full circle.

This circle happens in many places, both business and IT. Metadata management plays a role in that it supports the management of reusable elements of work (be they software components or financial instruments), once they are standardized. There is a spectrum of volatility: at the one end (immature, poorly understood, highly creative use) people should better be left to themselves than to put them in a straightjacket; at the other end (mature, thoroughly understood, highly mechanized use) it is almost imperative to leverage metadata so as to be flexible, confident, and fast at the same time.

The question is where to draw the line between creative and mechanized. Obviously, the answer has to do with the level of understanding of the objects in focus (components, instruments), their number (few, many, numerous), and the rate of change (quickly, slowly). Answering this question is one subject of this book: setting the peculiarities of a domain (financial services) in perspective with the needs of management (performance, risk, compliance) to arrive at a more refined understanding of metadata's value.

As you have seen, metadata management is not *l'art pour l'art*, a purpose in itself. It must create value. And exactly the question of value creation cannot be answered positively in all cases. In fact, sometimes it may destroy value, namely when it introduces unnecessary

bureaucracy to your business processes. All its splendor comes at a cost. Therefore, in situations where it is difficult to understand a problem that needs to be managed for risk (or compliance, for that matter) – for lack of knowledge, skill, time, because of political adversities, or whatever – there are factors that can take away a large part of metadata’s allure. It can become too costly, and you must balance such conflicts wisely. This book spends a great deal of time elaborating on the question of value creation and the conflicts that come with it.

### 1.3 Objectives

Having worked in the financial services industry for a while now, and having gone through the experience of designing, implementing, and operating a number of different metadata management solutions during that time, I cannot avoid but call it a mixed experience. I have long wondered about the reasons.

There were successes – some surprising, others more planned and expected. I still remember the evening I met with a former colleague for dinner before a night of cinema. He mentioned to me in passing that a credit product modeling solution I had designed before leaving the company had gone productive and been very much appreciated by business users. I was electrified: although we had been as thorough about the design as was possible, it did not seem to us there were too many people waiting desperately for our solution to be rolled out. Obviously we had struck a chord with enough of them, nevertheless. At the time of designing the solution, I was mostly concerned with a flexible architecture and integration with the software framework on top of which we had built it, but not with the business implications and possible impact on product design. I went home on cloud nine after cinema. The solution (in extended and evolved form) is being used to this day, six years after.

But there were also failures. The objective of this book is – to an extent – to demystify metadata management and explain some of this apparent randomness. In fact, it turns out that there is a host of good reasons why the projects I was involved in turned out the way they did. But this can only be understood when the analysis combines very different and, perhaps not surprisingly, mostly non-technical fields. The first aspect examines activities and their contribution to value creation so as to make them more economical. In the context of this book, we will take a particular look at the leverage effects of metadata management and how they contribute to value creation. The second area is industry peculiarities, i.e. the idiosyncrasies of banks and insurance companies. This sector exhibits its very own set of rules, based on its history, the nature of its value chains, and most notably, the role of national legislation. In particular, this book takes a closer look at some of the constraints that influence solution design.

In fact, there is a colleague of mine who likes to joke: ‘May all your problems be technical!’ On a positive note, there is a host of reasons to defend metadata management as an indispensable tool in managing complexity and volatility. Another objective of this book is hence to show



how the use of metadata can contribute to an effective alignment, leading to better change, risk, and compliance management. One reason is increasing systemic complexity, a burden that is attributable in part to the surge in the number and sophistication of regulations introduced in recent years, as well as the continuing evolution of products offered to the financial markets. This book will point out how metadata can critically enhance the ability of an institution to manage change, risk, and compliance by structuring the problem and actively managing transformation processes. Another reason is the widespread use of IT in financial services: banks and insurance companies are some of the most avid adopters of information technology. They have thus long streamlined many of their products, processes, and services (specifically in retail business) and achieved great productivity improvements. This in turn makes them perfectly suited to leverage techniques that build on top of those structures, which is at the heart of this book's idea: systematically managing and aligning the IT supporting the operation of a financial services company with its stated business goals, and continuing to do so in the face of change.

The problem of managing the complexity occurring in financial services is not a recent phenomenon. In fact, it has been around for plenty of years. Yet, many of the (sometimes excellent) publications on the matter have focused on but one side of the coin, either the business or the technology viewpoint. My experience is that this falls critically short of some of the problems you face in this area. Therefore, the overall objective of this book is to tread the line between business and IT, to provide you with a fresh view on things that combines insights from both areas. The most important of these ideas are explained next.

### 1.3.1 Value proposition for a vertical domain

Imagine a large bowl of spaghetti. Depending on where you poke your fork in, you may have a hard time extracting just enough spaghetti to fit into your mouth. The larger the bowl, the more likely you will end up with a huge pile of spaghetti on your fork, impossible to consume in one bite. However, if your host piles a smaller amount of spaghetti on your plate the whole exercise gets a lot easier. And as you may well know, the more you have eaten the easier it gets.

For a systematic understanding of our problem we must have a way of structuring it. When trying to understand the structure and rules underlying an area like financial services we need to do kind of the same: distribute spaghetti on plates, and adjust the amount of spaghetti to the size of the plates. This is the principle behind *divide and conquer*.

In IT, people often refer to a *domain* when talking about an area of knowledge or activity, like operating systems, relational databases, software development, or more business-related areas like risk underwriting, claims management, credits, payments, investment, or financial accounting. Domains are used to group the concepts used in the real world and their relationships into a coherent whole in order to manage them systematically, independent of others. This aids their management in just the same way that spaghetti plates help you find a group of spaghetti small enough to swallow.

There is a difference to spaghetti bowls, however. Domains are used to arrange and structure your view of the world. This structure often determines which relationships between the concepts you regard as interesting and which concept belongs where. By delineating domains in a specific way you emphasize certain aspects, while de-emphasizing others.

Traditionally, when the real world can be looked at from two fundamentally different viewpoints, people have distinguished so-called *horizontal* from *vertical* organization of domains. Putting things a little simply, a horizontal domain organization tries to maximize your own satisfaction, that is to put concepts into domains the way you prefer. On the other hand, vertical domain organization puts the other (possible) viewpoint first and arranges concepts in the way a person with that point of view would do.

Obviously 'horizontal' and 'vertical' are relative terms. Why all this complication? Value creation is best analyzed when arranged around the concepts dominant to the business process under study. Now, as one moves away from the purely technical (read: IT-related) world closer to business, it becomes less and less helpful to use technical concepts. Since the relationships between the concepts in a business domain are often different from those exhibited in a technical domain, the traces of a value chain can no longer be followed properly.

This book takes a vertical domain viewpoint. In my experience this is indispensable for understanding the levers of value creation at the disposal of business. On the flip side of the coin it helps tremendously in identifying areas where the use of metadata does not add value.

### 1.3.2 Tradeoffs in large multinational corporations

As was mentioned already, one focus of this book is on large, multinational, globalized companies. This is important for two reasons.

First, such firms are bound by the typical constraints of organizations beyond a certain size, which grew in a more or less uncontrolled fashion over an extended period of time. All of these constraints *decrease* our freedom in designing solutions:

- *Strong division of labor*: people in large companies tend to specialize. New techniques are not introduced by waving a magic wand. It takes time for people to understand what is intended.
- *Business heterogeneity*: there is a multitude of exceptions, niches, and variety in the business. When everything is special, abstracting from it becomes difficult.
- *IT heterogeneity*: banks and insurance companies adopted IT a long time ago and invested substantial amounts of money in their applications. These are not easily done away with and must be treated as monoliths, which does not always conform with plans. Hence, problems must be managed at very different levels of granularity.
- *Politics*: it is a matter of fact that quite a few decisions in an organization are not taken purely based on rational facts. Given the far-reaching consequences of leveraging techniques such

as metadata, political problems are to be expected. They are indeed so important that they must, in my opinion, be addressed strategically.

Second, large multinational corporations are particularly struck by the travails of globalization, which *increase* complexity:

- *Publicly held*: the way they are financed, and because of the critical role data plays in financial markets, means that public companies are the subject of much regulatory scrutiny and reporting requirements, much more than privately held firms.
- *Substantial size*: bigger multinationals are typically conglomerates of dozens, sometimes even hundreds of legal entities. They operate in their specific legal environment, which not only concerns regulatory or legal affairs, but also financial management issues like taxes.
- *Spanning jurisdictions*: legal entities, by definition, must fulfill certain legal obligations stipulated by the authorities. The financial services industry is among the most heavily regulated, which can dominate the way in which a firm's operations are structured: organizationally, in terms of products, or the way its business processes work.

The combination of forces influencing the design of a (successful) solution for a large organization is unique, and yet occurs often enough that it merits dedicated treatment. This is one objective of this book, namely to picture the tradeoffs effected by these forces, and outline a constructive way of dealing with them.

## 1.4 Scope

This is a conceptual book. It is not concerned with technology or implementation, at least not significantly. It is mostly concerned with business processes and data, and how both are supported, stored, and manipulated by applications, respectively.

Obviously, this book is about the financial services industry, banks, and insurance companies. You may find time and again that some statements would be applicable to other industries as well. I will not dwell on such similarities.

Geographically, or should I say, jurisdictionally, this book is international by design. It addresses the problems of multinational corporations, trying hard not to assume anything is common that in fact varies between countries. Many problems elaborated inside this book are driven genuinely by jurisdictional developments. The focus will be put on important financial markets like the United States, the European Union including the United Kingdom, Japan, and Switzerland.

This book covers both wholesale and retail business. Naturally, not every aspect can be covered exhaustively, so I will restrict myself to a sample of problems that can easily be

extrapolated to other areas, such as investment banking, reinsurance, primary insurance, or retail banking.

As a general rule, the book will cater to the needs of diversified companies, that is those with a wide-ranging, non-correlated portfolio of business activities. This does not mean that pure investment banks or specialty reinsurers are not part of the picture. In fact, you will later see that statements calling for specialization will be made in the interest of value creation. However, the point of view *generally* taken will be that of a large, diversified financial institution.

Likewise, the issues of risk and compliance management are a wide field. I would like to reiterate my statement that there are plenty of good books available. Furthermore, there is a plethora of regulations and frameworks that have been issued by national and international authorities, often based on laws agreed by the respective parliaments. Not all of them can be covered. Hence, this book will address selected pieces of regulation that have or can be expected to have a large impact.

Finally, value creation will be analyzed at a qualitative level only. The main reason for this is that coming up with quantifiable benefits depends wholly on the properties of the environment. This will be impossible to do, at a general level, in a book like this. Also, I believe that mathematical accuracy pales in comparison to the, well . . . qualitative differences in value creation that one can observe between, say the various lines of business in a financial institution.

This book takes a reactive perspective on change management. For the purpose of this text, it is always events elsewhere that trigger change processes. Initiating changes and managing them accordingly is a totally different field and will not be covered.

Of the specific types of risks, legal, political, and reputational risks will receive limited or no treatment. The main reason is that whereas the above carry at least some relevance for (operational) risk management, they require their very own ways of handling, which are not readily accessible to support by metadata. By the same token, risk management entails not only risk control, but also risk diversification and avoidance. These are out of the scope of this book as well. The obvious reason is that risk avoidance is a way of getting around a specific risk, at least in part. As will be shown later, the main focus will be on modeling things potentially imperiled by risk and trying to understand concrete perils and risk levels using metadata. Risk diversification (as in portfolio management), on the other hand plays an important role in managing investment and insurance risks, but has been extensively covered by publications elsewhere.

The value proposition, as has been said before, will not be covered at the quantitative level. By extension, covering the quantitative aspects of risk management (during risk assessment) would bloat the contents of the book as well. It would also probably not add too much insight to material presented at the conceptual level, anyway. Again, please turn to appropriate literature elsewhere.

Many of the recommendations that are to follow are based on a model in which metadata management is carefully planned, executed, and monitored. Especially when it is used to support understanding (e.g., as in descriptions of the business meaning of data or processes), there is at least one other model that can be used, namely so-called *social tagging*, which is a community-based, cooperative mechanism for classifying and documenting things. Typical examples of social tagging include the photo sharing platform Flickr and the online (book) retailer Amazon. Both offer their users ways of adding pieces of data to objects served on their site, such as book ratings or comments on Amazon, or descriptive keywords on Flickr. There are two reasons for not covering this model of managing metadata:

1. The dynamics of social, community-based interactions are only slowly being understood, so it would be a little early to cover them here. Furthermore, it is difficult to scale them since they are based on cultural homogeneity and decentralization, whereas this book focuses on large, heterogeneous organizations and central control.
2. Value in social tagging is (also) created through network effects, the use of which has a completely different underlying economic model than the two main effects presented here, namely productivity and quality gains.

It should be mentioned that there are situations where social tagging is preferable to the planned model. For the sake of simplicity such scenarios are not considered. For all practical purposes, social tagging should be considered complementary to planned approaches, and as such you may add these mechanisms on your own where you find them helpful.

## 1.5 Who should read this book?

This book will appeal to three groups of people, corresponding to three organizational functions:

- *Information technology*, organizing the furtherance of applications for the benefit of the corporation, specifically in the form of an integrated framework to align its capabilities with strategic business goals.
- *Risk management*, setting standards on the systematic identification, evaluation, acceptance, reduction, avoidance, and mitigation of risks in which the corporation engages, as well as enacting effective controls to ensure compliance.
- *Compliance management*, responsible for ensuring the conformity of operational business activities with internally formulated or externally given rules, specifically with regards to regulatory authorities.

In terms of organizational level, this book appeals to the first two or three lines:

- *Chief officers*, managers representing one of the above functions at the executive board level, taking decisions of strategic impact.
- *Senior executives*, direct reports of chief officers, responsible for running the day-to-day operation of select parts of an organizational function.
- *Specialty managers*, people who have profound expertise in a specific area, typically serving in a consulting role to senior executives and chief officers.

This book primarily intends to address the following points of interest to the above groups:

- Chief Information Officers (CIOs) and senior IT executives learn how to achieve change goals and track progress systematically. They will understand where and how metadata management ties in with IT processes, what the levers of value creation are, and how they relate to overall value creation in a financial services organization. Furthermore, they get a picture of the downsides and practical problems and how to mitigate them. Finally, they learn how to start a comprehensive process towards metadata management.
- Chief Risk Officers (CROs) and senior risk managers learn the nature of metadata and its use in structuring and modeling risks. Specifically, they learn how the management of operational IT risk can be systematically improved by establishing and maintaining risk portfolios described by metadata. In addition they learn how metadata helps them gain better insight into risk data in general. They will also understand the impact on the organization, applications, and operative processes, as well as the pitfalls.
- Chief Compliance Officers (CCOs) and their senior managers learn how governance, compliance and metadata are related. They formulate and adopt regulations in a methodological fashion. Specifically, they will understand how regulatory alignment can be established, properly assessed, and re-established once deviations are discovered using architectural alignment of IT with stated rules. They will also understand the limitations of such an approach, the costs incurred, and practical mitigation measures.
- Chief Business Engineers (CBEs) and business architects learn their supporting role, the impact of a particular way of scoping and designing architecture, domain model, and metamodel, and how this impact affects the organization's ability to manage change, risk, and compliance actively. They also learn how they themselves can use metadata to improve the results of their own work. This especially applies to capturing and analyzing existing IT portfolios (*as is*) and setting targets for their evolution (*to be*). Finally, they manage to limit the adverse effects on the work of projects, developers, and business users.
- Product managers learn to better understand how to leverage the power of metadata in designing, deploying, and changing products in a complex and continuously evolving environment. They will understand the prerequisites that IT systems must fulfill in order to

support them, but also the organizational setup required and skills needed. Furthermore, they learn how to reduce complexity in the right places, and yet maintain the ability to handle it in others.

Obviously, people working in teams directly supporting the above groups will also find the book useful for the same reasons. As a kind of secondary goal of this book, the following groups of people will find reading the material helpful:

- Business analysts will understand how metadata affects their work, what is asked from them to support it, and what criteria are important in this.
- Methodologists will learn how the processes of analysis, design, construction, and operation are affected and can be supported by metadata to improve productivity and quality.
- Expert business users will see how they can improve their efficiency and effectiveness when faced with complex tasks. They specifically will understand metadata quality deficits and how to handle them.

Finally, there are the specialists in all the different domains (programmers, underwriters, etc.). I do not list them here, because they are typically quite detached from much of the above, but they may use the material for information purposes.

## 1.6 Contents and organization

This book is divided into three parts. First, it illustrates the peculiarities of metadata management and its use in architectural alignment. Second, it outlines how and where metadata management creates value, namely by supporting an efficient response to change, improving the ability to ensure compliance and regulatory alignment, and supporting the management of risk. Third, it casts a light on the practical problems (the mechanics, if you wish): dealing with business and technical evolution in its various forms, managing quality, and planning, extending, and sustaining the success of a metadata-based solution.

Chapter 2 takes a reactive, process-oriented perspective on the management of change in a corporation. Business processes are structured into core, support, and steering processes. Metadata management is presented as an IT support process controlling and supporting the adoption of change in a company's IT. The two main facets of metadata management are described: metadata as a reification of how models are related, and metadata management as a process that couples change adoption activities. With this, you achieve a separation of concerns at the data level, while requiring an integration of concerns at the process level.

Chapter 3 presents alignment as an activity that uses metadata by reifying how compliance is achieved, and then actively managing that metadata as part of governance activity. The chapter

emphasizes the interplay of architecture and metadata, describing how architectural building blocks relate to metamodel abstractions, and how that can be used to ensure compliance. It is highlighted how metadata can be made a strategic tool in the planning and implementation of structural changes. Finally, the important practical issue of handling exceptions to the rules is given room, with the text explaining how to deal with them from a data and process perspective.

Chapter 4 concerns itself with making the organization more efficient and effective in responding to change, in a word: to achieve productivity improvements. The different cost factors arising through the use of metadata management are discussed and juxtaposed with the realities of the financial service industry. The effects of typical changes on the IT landscape of companies are discussed, leading to the result that only a targeted use of metadata management creates significant value. The main message is that a bank or insurance company should seek performance improvements through metadata management only in its company-specific way of adopting technology. Several case studies illustrate how others have used metadata in this way, such as credit product design, using business rules in trade automation, and model generation in asset management.

Chapter 5 looks at the risk management process and describes how you can make use of metadata to manage risks more successfully. First off, it explains why IT today plays a bigger role in managing risk than ever before. It then goes on to illustrate what effects the use of IT has in managing risk, both direct and indirect. It also emphasizes the risks IT itself introduces. With this picture in mind, the chapter claims that the role of metadata in managing risk lies in the support for risk transformation, data quality management, and architectural alignment. Several case studies illustrate the practical use of metadata, ranging from data quality management in data warehousing to operational risk management in architectural governance.

Chapter 6 gives a brief overview of what compliance means in terms of regulation, governance, control, and audit. Then it takes a look at the regulatory landscape in different jurisdictions, and discusses the nature of different regulations. From there, it explains what role IT can play in supporting compliance management. The contribution of metadata management to this is explained. A number of case studies from anti-money-laundering, SOX- and US-GAAP-compliance, taxes, privacy, and documenting the opening of relationships complete the picture.


The first practical issue is handling evolution, which is at the heart of Chapter 7. The activities in this area are listed and explained, such as metamodeling, organizing impact analysis, or change propagation. In all this, the challenges of real-life large organizations and how to overcome them are woven in, such as dealing with division of labor, dealing with impedance, retaining and regaining consistency, or grouping changes. The various tradeoffs of solution templates are discussed.

Chapter 8 looks at how metadata quality deficits affect its utility. Because metadata management is a support process, any deficit in its intelligibility, completeness, correctness, consistency, actuality, or granularity has a characteristic impact on the core processes it ties



into. Ways to detect such deficits and countermeasures for tackling them are explained from a risk management perspective. Side effects of such countermeasures are explained.

Chapter 9 deals with the question of how to ensure the continued success of metadata management in typical big corporations, despite such usual phenomena as politics, lack of ownership, or limited awareness and understanding. Ways of identifying and recruiting owners, as well as alleviating their fears are listed, also with respect to other stakeholders. The chapter also discusses the challenges in scaling metadata management to global scope. Awareness and understanding, the issue of training people, and making them use metadata appropriately, are problems that receive coverage as well. Finally, the issue of company politics is examined, its nature illustrated, and the reason why metadata management gets so tied up with it is explained. Different political maneuvers are discussed, highlighting both the offense and defense of playing politics.



**Figure I.1:** Context diagrams feature a single process, which sets the scope of attention. Outside actors, such as roles or organizations with which the processes interact, are grouped around it

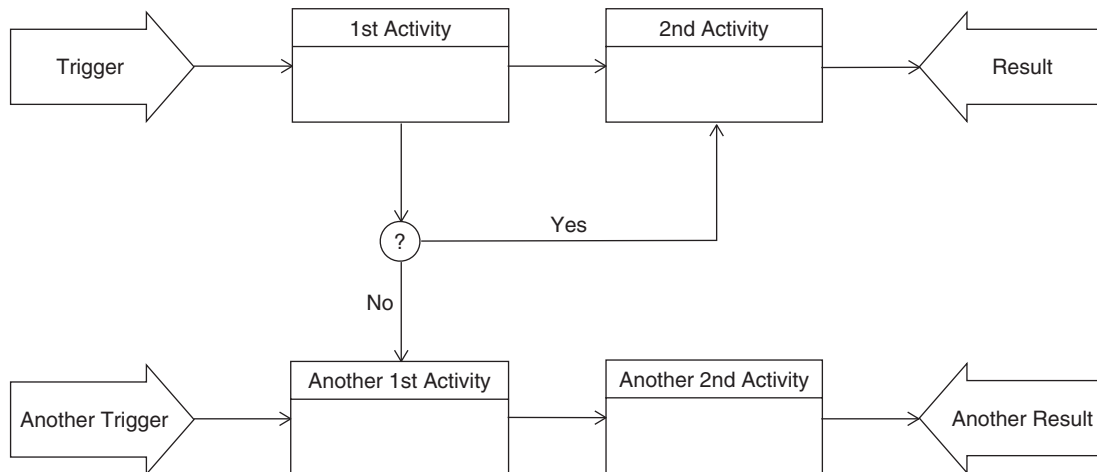
Finally, Chapter 10 puts forth a vision of an institutionalized corporate function for combined change, risk, and compliance management. It revolves around using models for a more systematic management of change, using the mechanisms of enterprise architecture and metadata management in union. The tasks of such a corporate function are described. In closing, the main lessons (some dos and don'ts) learned from this book are summarized, and a forward-looking statement concludes the book.

The notation used for diagrams is as follows. The *context diagram* (see Figure 1.1 for an example) is used to illustrate what environment a business process is embedded in, and how it interacts with that environment. It sets the scope of the discussion by defining what is considered part of the problem and what not.

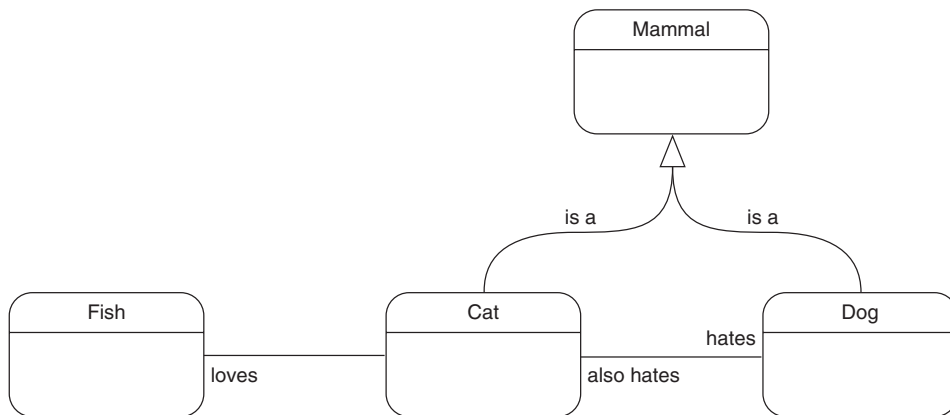
The *process diagram* (Figure 1.2) describes the events triggering activities in processes, the sequencing of activities, and their completion in the form of result events. Process diagrams specify how processes are executed.

The *data diagram* (Figure 1.3) is an entity-relationship model that illustrates the main data elements and their relationships at a high level. Data diagrams are not complete, but merely serve to highlight the most important connections.

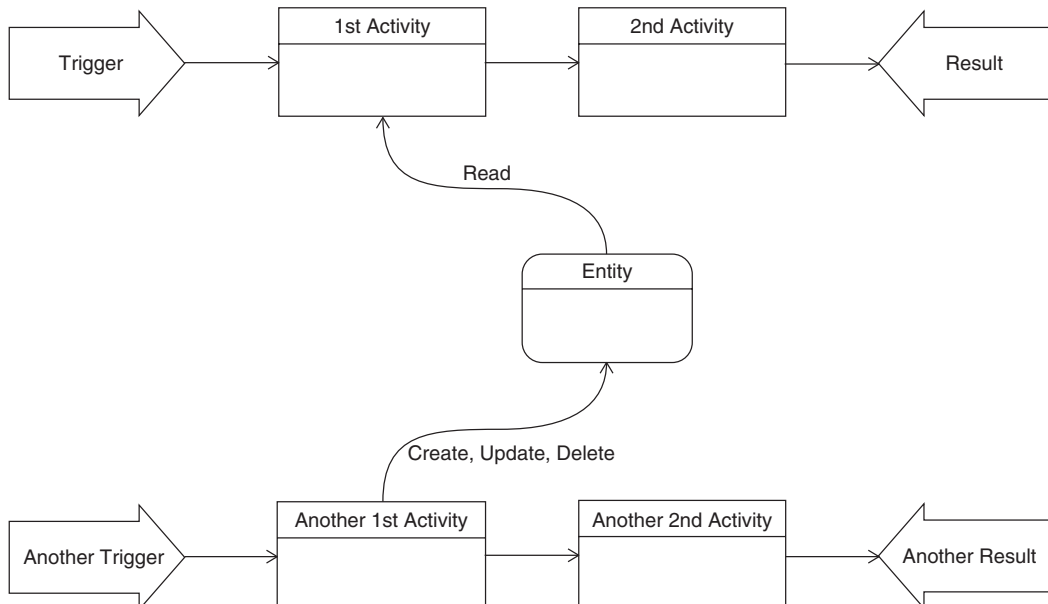
The *data flow diagram* connects processes and data (Figure 1.4), and highlights the most important interactions of two independent processes via the data that they exchange. A data flow diagram can be used to illustrate what data originates where, and what route it takes between processes. Its function is to describe the sequence of interactions.



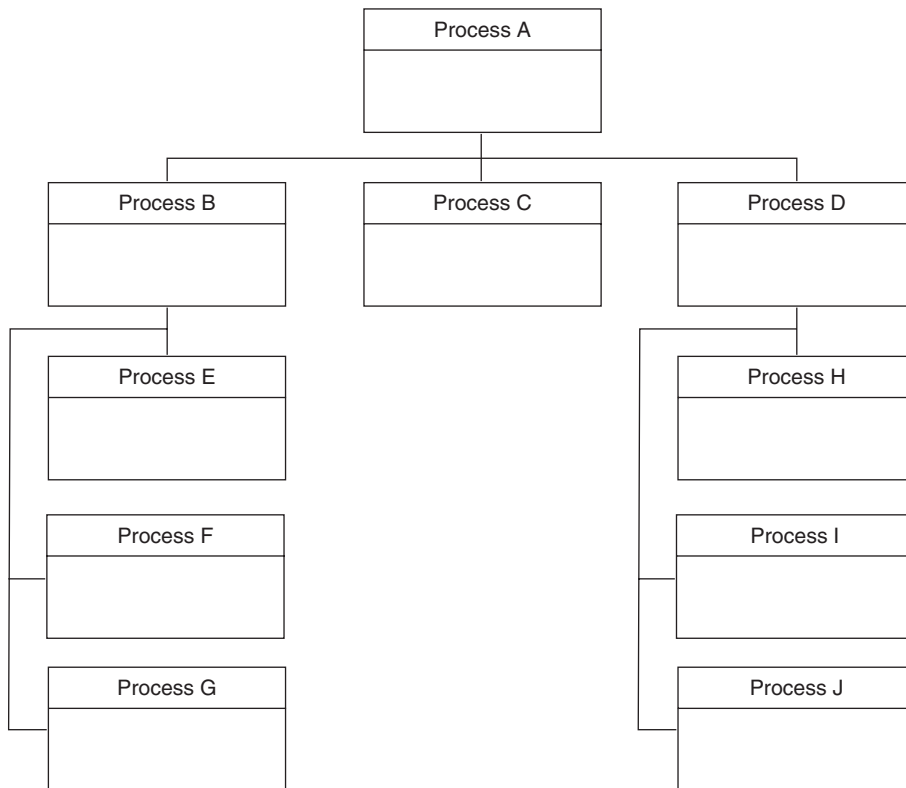
**Figure 1.2:** Process diagrams illustrate the events triggering activities and the results thereof. Branching is used where necessary, with the labels on the branch describing the choice



**Figure 1.3:** Data diagrams describe the structure of the data being used, that is entities and their relationships. Where necessary the relationships are labeled to describe their meaning. The triangle signifies the 'is a' relationship, meaning that the entities subsumed under it are treated the same way by processes as the entity above



**Figure 1.4:** Data flow diagrams describe the interaction between data and (at least two) processes. Each process shows an end-to-end activity sequence from trigger to result and which entities it uses (reads) or manipulates (creates, updates, deletes). The direction of arrows thereby indicates which way data is being used. Later diagrams will not label arrows. Data diagrams are used to illustrate how processes are coupled to each other through the data they exchange



**Figure 1.5:** Process hierarchies illustrate which processes and activities belong together. A process in a lower layer ‘belongs to’ the one above, meaning that the subordinate is primarily executed as part of the superior one

Finally, the *process hierarchy* specifies which process is part of another process (Figure 1.5). Typically, activities can be subdivided along some criterion, and the process hierarchy is there to illustrate that subdivision. With it, you can see which activities are part of others.