RZ 3656 (# 99676) 08/21/2006
Computer Science 25 pages

Research Report

Dynamic and Risk-based Compliance Management

S. Müller and C. Supatgiat

IBM Research GmbH
Zurich Research Laboratory
8803 Rüschlikon
Switzerland

Dynamic and Risk-based Compliance Management
S. Müller and C. Supatgiat

The necessity of complying with an evolving set of regulatory requirements is a growing concern to enterprises. Responsible decision makers must continuously decide which measures are appropriate and must be implemented with which priority in order to reach the optimal compliance level. To proactively address external audits, management must also decide on the optimal way to internally inspect whether the taken measures are effective and being followed. In this paper, we propose a quantitative risk-based compliance management approach, which allows management to optimally and dynamically select feasible measures to attain an adequate compliance level and to inspect compliance with a given set of regulatory requirements. We strive to minimize the expected total cost of compliance including the costs of individual measures and inspections, and the audit outcome cost for varying compliance levels. Our approach is based on dynamic programming and naturally accounts for the dynamic evolution of the enterprise with respect to the regulatory landscape governing it. The main merit of our method lies in its use as a scenario-based management support system. Depending on the availability and accuracy of input data, it can even be used as a comprehensive tool to optimally select the desired compliance measures and controls policies. Moreover, our tool lends itself as a policy instrument and may provide valuable guidance to effective rule making.

1. Introduction

Following a number of corporate accounting scandals, privacy breaches, money-laundering activities, and the raising threat of international terrorism, recent years have brought a surge of new and rapidly evolving regulations and provisions imposed on enterprises. Affected enterprises are challenged to continuously adapt their operations to these new requirements and to periodically demonstrate compliance with relevant regulations. In addition, regulations such as the prominent Sarbanes-Oxley Act (1) or the USA Patriot Act (2) have significantly heightened the expected costs of non-compliance by increasing the maximum penalty (i.e., increased fines and the possibility of imprisonment) and by making corporate directors directly liable for neglecting to properly address these regulations. As a result, large and listed corporations are currently spending enormous amounts of money in their attempts to achieve maximum compliance with the relevant provisions. The costs involved are enormous and cannot be blindly justified to the shareholders.

While attempting to attain perfect compliance is a noble goal, large enterprises are generally too complex in nature as to allow responsible management to know each and every detail that might possibly affect the companies' compliance risk exposure. Compliance management and the involved decisions on the targeted degree of compliance and the prioritization of compliance activities are inherently risk-based. To ensure a near-perfect degree of compliance, a company would have to employ a large number of internal auditors to inspect their employees, systems, processes, and products

on a daily basis. While such a compliance management strategy would absorb an enormous amount of financial resources, it would still fall short of providing perfect compliance with certainty. Hence, compliance is a continuous rather than a binary phenomenon and must be measured on a ratio scale to allow for sensible decision making. Accordingly, the amount of effort an enterprise puts into achieving compliance must be a direct function of the scrutiny, frequency, and outcomes of conducted audits, the cost and effectiveness of possible measures to implement compliance, and the likelihood and impact of implicit and explicit costs for varying degrees of compliance with a particular regulation.

Like any other business decision also compliance management needs to be conducted in a prudent and proactive manner, that is, also compliance-related activities must be carefully analyzed and prioritized with respect to their potential benefits and costs and they must be financed using scarce resources. Because of the scarcity of resources (time, money, people, etc.), the implementation of measures to comply with regulatory requirements needs to be prioritized according to their expected cost. We show how this can be done. All available measures to implement some regulatory requirement provide at least an adequate compliance level, and differ only in their cost, implementation time and the amount of more-than-adequate compliance that they provide. And given that the compliance level of a company decays over time and that various audit outcomes lead to different costs of respective auditor recommendations, our model finds the *optimal portfolio of measures* to be implemented over time.

In our model, we assume that there is always the possibility that some people or systems do not respect the legal obligations. The goal of our paper is to show how to manage such risk. Many factors influence both an enterprise's compliance degree and the implied compliance risk (i.e., the expected cost of compliance). Among them are the types, effectiveness and cost of possible measures to address a specific regulatory requirement, the type and frequency of inspections conduced, and also the audit coverage and audit outcome costs.

The expected costs and benefits of all possible compliance measures must be analyzed with respect to a number of uncertain future outcomes. While risk management has a long-standing tradition in areas such as finance and insurance to manage financial risks (3,4,5), credit risk (6), and recently also operational risk (7,8,9) and IT security risk (10,11), to our best knowledge, we are the first to address the management of compliance using a truly risk-based approach. Ironically, while financial services companies have made risk management to one of their core competencies, they do not seem to have realized that also the unpredictability of regulatory change should be addressed using similar techniques.

In this paper, we introduce a dynamic and risk-based approach to compliance management. Our method treats the recurring decisions on what measures to implement to achieve compliance, and how to test their effectiveness as inherently risk-based. As a result, an enterprise, or organizational unit, employing our approach, will manage its targeted compliance risk level by taking into account the cost and effectiveness of

previously implemented and future measures, the type and cost of internal inspection, the auditor's scrutiny and effectiveness, and the likelihood and expected cost of audit outcomes. Our approach can be implemented in a straightforward way and be used as a scenario-based management support system in order to implement the optimal portfolio of measures to maintain the desired target compliance level. Our model can also be used as a policy instrument and support effective rule making.

Measures and Inspections

When attempting to attain compliance with a given regulatory requirement, on the one hand, enterprises need to select the most appropriate *measures* to achieve the targeted degree of compliance. On the other hand, they also need to decide on a suitable way to *inspect* whether the chosen measures work as desired. Inspecting the deployed measures informs on the current compliance status and generally leads to an improved compliance level of the enterprise. This also means a higher likelihood of passing a future audit with a higher satisfaction level and avoiding the (implicit) cost related to passing the audit with lower satisfaction or, in the worst case, failing it. Measures and inspections to address compliance concerns often entail significant costs. Specifically, compliance-related investments represent opportunity costs, in that they bind money that could otherwise be used for other, more urgent or potentially more lucrative investments. Furthermore, different types of measures and inspections are not equally effective. That is, depending on the actual set of measures chosen, the compliance level of the enterprise can be increased more or less contingent on the relative effectiveness of the selected measures.

Different kinds of regulations necessitate a variety of measures and inspection types. Most regulations have proprietary terms to refer to what we abstractly denote as measures and inspections. For example, in the context of Sarbanes-Oxley compliance, the law recommends that affected enterprises implement the control framework of the Committee of Sponsoring Organizations of the Treadway Commission (COSO) (12) in order to protect the accuracy of financial data and to become and remain compliant with relevant laws and regulations. Towards this end, COSO also defines two concepts, which represent the notion of measure and inspection: controls and their testing. Other regulations know similar concepts. For instance, privacy regulations usually speak of appropriate access control measures that need to be deployed to protect personal data. Often, the effectiveness of such controls also needs to be tested in order to credibly demonstrate that the relevant privacy provisions are being followed, thus representing a particular kind of inspection.

According to COSO, a control is "a process, effected by an entity's board of directors, management, or other personnel, designed to provide reasonable assurance regarding the achievement of objectives in the following categories:

- Effectiveness and efficiency of operations
- Reliability of financial reporting
- Compliance with applicable laws and regulations"

Furthermore, COSO stresses the important role of people at every level of an organization to become and remain compliant. It also recognizes that "internal control can be expected to provide only reasonable assurance, not absolute assurance, to an entity's management and board."

This characterization of the notion of control realizes the risk-based nature of compliance-related measures by defining them as providing merely relative and not absolute assurance. Furthermore, it is recognized that compliance is an ongoing concern and, hence, must be continuously monitored and adapted.

External Audits and the Costs of Compliance

In some cases, regulations provide direct economic value to affected enterprises (as in the case of e.g. privacy regulation). In other cases, regulations constrain the scope for conducting business, thereby inducing direct and opportunity costs on enterprises and requiring them to implement costly compliance measures. In such cases, the incentive of implementing appropriate measures and conducting inspections on their effectiveness with respect to other regulations would not be very large if it were not for external audits that periodically occur. An (external) audit is conducted by a number of auditors and denotes the evaluation of an enterprise, specifically, of its systems, processes, or products. The purpose of an audit is to verify that an enterprise operates according to a set of relevant regulatory requirements. There are different types of audits (e.g., financial audit). Some audits are also done voluntarily by the enterprises to prove their conformance with certain standards and become certified (e.g., ISO) or to get access to a certain market. We focus on audits of regulations in this paper.

Audits can be passed with varying degrees of auditor satisfaction. Depending on the outcome of an audit, an enterprise may have to implement different amounts of recommendations. In general, such recommendations lead to the implementation of additional or stronger measures. This constitutes additional cost to the enterprise. If the enterprise does not implement measures that provide an adequate compliance level, the enterprise may fail the audit and face a very high penalty.

Passing an audit with high auditor satisfaction generally implies that the enterprise in question must only take minor corrective action, whereby it incurs a certain cost. Normally, the audited company is granted a certain grace period for the correction and adaptation of the reprimanded system, process, or product. However, there will be no additional audit before the next regular audit period.

If an enterprise passes an audit with a low level of auditor satisfaction, in addition to requiring that the identified deficiencies be corrected, a governmental agency, informed by auditors, may impose a certain fine on the enterprise. Both the correction of identified deficiencies and the potential fine represent explicit costs to the enterprise. After the audit, the enterprise is also granted a grace period within which is must be able to attain and demonstrate a satisfactory compliance level. Subsequently, the enterprise may also be

required to pass a follow-up audit, wherein the auditors investigate whether the identified deficiencies have been corrected.

Abstractly speaking, the outcome of an audit depends on two dimensions: the compliance level of the enterprise and the audit coverage or scrutiny of the auditors. If the enterprise has invested heavily in measures to attain a high degree of compliance with relevant regulations, even broad audit coverage is not very likely to reveal a state of inadequate compliance. However, if the overall compliance level is relatively low, broad audit coverage (i.e., a high probability of detection) will likely translate in revealing that the enterprise is only minimally compliant and passes the audit with lower satisfaction, or in the worst case, is not adequately compliant and thus fails the audit. In this paper, we adopt the convention that we only talk about a failed audit if the auditors found that the measures implemented by the enterprise did not lead to an adequate compliance level.

Over time, enterprises accumulate a historical audit track record that influences the behavior of its auditors. Auditors are less likely to apply broad audit coverage when auditing an enterprise with an outstanding audit track record (i.e., which has passed the previous k audits with high auditor satisfaction). In contrast, auditors generally apply more scrutiny when they are confronted with a company that has shown a lower compliance level in the past.

The expected cost of the different compliance levels includes explicit and implicit costs:

Explicit costs include:

- Implementation costs for implementing auditor recommendations.
- Prohibition of business expansion activities during the time adequate compliance is being regained following audit outcomes with low satisfaction only. This implies missed opportunity of business growth and represents a high (opportunity) cost for most companies, measured in lost revenue, lost market share and reduced competitiveness.
- Prohibition of selling a specific product (e.g., drug, etc.) while compliance of the product in question is not satisfactory.
- Monetary fines imposed on the enterprise.
- Personal and criminal liability for CEOs, CFOs, auditors and board members for financial discrepancies and/or operating the enterprise in a state of non-compliance.

Implicit penalties include:

- Demand decline because of negative publicity (loss of reputation, bad image, and decreased customer respect). This may lead to a lower share price.
- Demand decline because of lack of trust in safety of product(s) with the result of a lower share price.
- Higher share price volatility because of insecurity of enterprise's future performance.

Not all auditor recommendations are equally expensive to implement. Likewise, not every regulation imposes equally severe consequences in case of inadequate compliance. As a result, different compliance levels with respect to different regulations translate to different expected compliance cost. As all resources of an enterprise (time, money, people, etc.) are limited, this is important to prudently prioritize compliance activities. Also the scrutiny of auditors, the enterprise's own compliance level and its audit track record are important factors when contemplating the expected cost of possible compliance activities and audit outcomes.

Structure

The remainder of this paper is structured as follows: In the next chapter, we will introduce a mathematical model for dynamic and risk-based compliance management. The model leads to the formulation of a total expected cost-to-go function, which will be minimized using dynamic programming. Subsequently, we will investigate a case study and demonstrate how the model and the respective dynamic programming algorithm can be used to implement the optimal portfolio of measures, inspection type and frequency that minimize the expected cost of compliance.

2. Model

The compliance manager dynamically manages the compliance risk with respect to a particular set of regulatory requirements over time. Let τ be the total number of time periods in the decision-making horizon. At the beginning of each period, the manager performs a risk assessment and decides on which measures to implement and which type of inspections to conduct with which frequency in order to address the previously identified compliance risks.

Decisions

There are M possible types of measures to address the regulations. Each measure has two cost components: a one-time fixed cost and periodic maintenance cost. Measure i costs c_i^M to implement and an additional y_i^M to maintain in each period. If there is no maintenance cost then y_i^M is set to zero. When implementing a measure i, it will take r_i^M implementation periods before the measure is successfully implemented and effective. The manager can also stop maintaining any measure she implemented earlier. The measure that is not maintained anymore will loose its effect in the period after the manager stops paying the maintenance cost. At a later point in time, if the manager wants to implement the measure again, she has to pay the full cost again as if the measure has never been implemented. Implementation will again take r_i^M periods.

There are I possible types of inspections with different costs and effectiveness. Inspection j costs c_j^I per inspection. Without loss of generality, in each period, only one inspection type is allowed. If inspection types a and b can be conducted in the same period, we can just define a new inspection type, say c, representing the combined cost and effect of both inspection types a and b. Therefore, conducting inspection type c yields the same effect as conducting both inspection types a and b.

We let an integer vector $V_t = [v_1, \ldots, v_M]$ be the historical measure implementation vector. Its i^{th} component represents the number of periods from period t until measure i will be effective. If measure i is already in effect, then $v_i = 0$. If measure i has never been implemented, then $v_i = -1$.

Regulatory requirements and Measures classification

There are J requirements to fulfill. Accordingly, the measures can be classified into J classes. Each class corresponds to one of the requirements that the measures aim to address. We represent the effectiveness of measure i to address requirement j by e^j_i. The effectiveness has a value between 0 and 1, with 0 denoting no effect and 1 encoding perfect effectiveness.

Compliance level

The major component in compliance level modeling is a *target compliance level*, denoted by $T(V_t)$. It is defined as

$$T(V_t) = \sum_{j=1}^{J} v_j \underbrace{\left(1 - \underbrace{\prod_{l=1}^{J} \underbrace{\left(1 - \max_{i \mid (V_t)_i = 0 \wedge i \in \Theta_l} \left\{ e_i^j \right\} \right)}_{NCSC}}_{NCR} \right)}_{CR},$$

where the symbol $(V)_i$ represents the i^{th} element of vector V and the set Θ_l represents the class of measures mainly addressing requirement l.

The target compliance level represents the maximally achievable compliance level given all measures that are currently in effect. It is computed as a weighted average of the individual compliance levels with respect to the J requirements. The weights v_j are assigned according to how much each requirement contributes to the total regulatory exposure. They should sum to one (i.e., $\sum_{j=1}^{J} v_j = 1$).

In the above formula, the term NCSC represents the minimal degree of non-compliance attained by implementing a set of measures from the same class addressing the j^{th} requirement. The term NCR represents the total non-compliance level resulting from implementing measures from different classes. Finally, the term CR yields the total degree of compliance with respect to the j^{th} requirement.

From the target compliance level formula, we see that if two or more measures from the same class are implemented together, only the one with the higher effectiveness for the corresponding requirement will affect the target compliance level. If measures belonging to different classes are implemented together, their combined effectiveness will define the target compliance level. The target compliance level is the maximal compliance level that can be obtained given the set of implemented measures. To increase the target compliance level, more measures or measures with higher effectiveness must be implemented.

We denote the compliance level of a company with respect to a particular set of regulatory requirements at the beginning of period t by a number b_t. The compliance level b_t is an indication of the company's current internal compliance level with respect to the set of relevant regulatory requirements. It takes a value between 0 and 1, where $b_t = 1$ denotes the highest compliance level.

If there is no inspection in period t, the compliance level at the beginning of period $t+1$ is a function of (a) the measures taken and the numbers of periods before they will be in effect, V_t and V_{t+1}, at the beginning of period t and $t+1$, and (b) the compliance level in the last period b_t.

When an inspection is conducted in period t, the compliance level is normally increased. The inspection effectiveness varies depending on the type of inspection chosen. Let a_t^I be an integer from 0 to I representing which inspection type is conducted in period t. The

value of 0 means no inspection is conducted. For inspection type i, the improvement is denoted by O_i, which is a factor, ranging from 0 to 1, denoting the increase in the compliance level b_{t+1}. A value of 1 means a full improvement with respect to the original level achievable by the implemented measures while a value of 0 stands for no improvement in the compliance level. The improvement O_i is assumed to be a random variable with probability distribution $F^{O}{}_{i}$. In any case, no matter how effective of the inspection, the maximum compliance level after any inspection is limited to the target compliance level corresponding to the measures being in effect.

The compliance level can decrease over time, for example, because the employees become more relaxed over time and do not adhere to the implemented measures so much anymore. We define the decay factor ρ, with a value between 0 and 1, as a multiplier to the current period compliance level to get the next period compliance level. The higher the decay factor, the faster the compliance level drops.

The compliance level b_{t+1} is a function f of $(V_t, V_{t+1}, a_t^I, b_t)$ and is defined as follows:

$$b_{t+1} \equiv f(V_t, V_{t+1}, a_t^I, b_t) = \begin{cases} \dfrac{b_t \rho T(V_{t+1})}{T(V_t)} & \text{if } a_t^I = 0, \\ \dfrac{b_t \rho T(V_{t+1})}{T(V_t)} + O_i\left(T(V_{t+1}) - \dfrac{b_t \rho T(V_{t+1})}{T(V_t)}\right) & \text{otherwise.} \end{cases}$$

The first case is when there is no inspection in period t. If there is no change to the effective measures, i.e., $V_t = V_{t+1}$, the next period compliance level b_{t+1} is just $b_t\rho$, i.e., the current compliance level with one period decay. If there is a change in the effective measures, then the new compliance level is equal to the ratio $b_t \rho / T(V_t)$ of the new target level $T(V_{t+1})$.

If there is an inspection in period t, the compliance level is improved by the amount $O_i\left(T(V_{t+1}) - b_t \rho T(V_{t+1}) / T(V_t)\right)$. Note that the term $\left(T(V_{t+1}) - b_t \rho T(V_{t+1}) / T(V_t)\right)$ represents the gap between the target compliance level and the actual compliance level under no inspection. The improvement factor O_i is multiplied with this gap to determine the improvement in the compliance level due to inspection i. If O_i is 100%, then the compliance level will be equal to the target compliance level. If O_i is 0%, then there is no improvement. Since O_i is a random variable, the compliance level b_{t+1} is also a random variable.

Auditing

We assume that auditing takes place every fixed interval. The inter-auditing interval is denoted by T_A. For example, when the time period represents one week and auditing occurs twice per year, then auditing occurs every 26 periods or $T_A = 26$.

There are N possible outcomes of an audit, ranging from outcome 1, i.e. passed with 100% satisfaction, to outcome N, i.e., failed with 0% satisfaction. We let K be the number of past audit outcomes sufficient to determine the audit outcome cost. That is, it is sufficient to calculate the audit outcome cost if we know only the past K audit outcomes. We let an integer vector $H_t = [h_1,\ldots, h_K]$ be the historical audit outcomes vector. Its i^{th} component represents the audit outcome, which is a number from 1 to N, at the i^{th}-last audit since period t.

The audit outcome cost at time t is a function of the historical outcomes vector. For example, the cost will be high after a series of consecutive bad audit outcomes. On the other hand, it will be low if a number of previous audits were passed with high auditor satisfaction. A company with a good auditing track record may only risk a warning or incur low costs if it happens to pass an audit with lower satisfaction whereas a company with a poor track record will incur a significant cost (e.g., due to the implementation of many auditor recommendations.). Furthermore, a high number of consecutively bad audits will also lead to more auditor scrutiny, yielding a higher probability of detection. If an audit is conducted in period t, the actual auditing coverage is denoted by q_t. The value of q_t is between 0 and 1, where $q_t = 1$ means 100% coverage. We model it as a function of the historical audit outcomes H_t.

$$q_t = g(H_t)$$

It is sensible to assume that audit coverage q_t will be high when the past audit outcomes were poor and will be lower when the past audit outcomes were good. This is because the auditors tend to put extra focus on companies with poor records.

The probability of detection from an audit depends directly on the audit coverage and the compliance level of the company. Broader audit coverage is associated with a higher detection probability. A lower compliance level implicitly reflects a higher number of less-compliant parts (i.e., components, systems, or processes) of the company or a moderate number of highly non-compliant parts. Hence, a lower compliance level is assumed to be associated with a higher detection probability. We have that

$$P(\text{detect when audit in period } t) = q_t\,(1\text{-}b_t)$$

In line with our definition above, auditors may reveal N-1 possible non-compliant states of the enterprise. Given the current compliance level b_t, the current audit outcome h_0 can still be uncertain and depend on uncontrollable factors outside the model. It is a random function of the current compliance level, i.e.

$$h_0 = U(b_t),$$

where U is a random function with distribution F^U.

Let d_t define the audit outcome cost incurred after the compliance level is audited in period t. We assume that the cost d_t is the result of a function z mapping the current audit outcome h_0 and historical audit outcomes H_t to a positive real number. That is,

$$d_t = z(h_0, H_t),$$

We shall assume that the cost d_t be higher for a worse audit outcome h_0. Furthermore, it is also possible that the auditors impose additional penalties onto companies with poor track records. In the next subsection we formulate our multi-period decision problem as a dynamic programming model (16)(17).

Dynamic programming model

There are three types of uncertainties in our model: the uncertainty of the inspection effectiveness (F^O), the uncertainty of detecting non-compliant behavior in an audit ($P_{(\text{detect when audit in period } t)}$), and the uncertainty of the auditing outcome after a non-compliant event is detected (F^U). At the beginning of each period, the compliance manager decides on which measures to implement and which type of inspections to conduct. We denote the actions in period t by A_t^M and a_t^I . The vector A_t^M is a binary vector of M elements, with its i^{th} element representing whether measure i is implemented or maintained in period t. The value of 1 means it is implemented or maintained in period t, while 0 represents a measure that is not implemented. Action a_t^I is an integer from 0 to I representing which inspection type is conducted in period t. The value of 0 means no inspection is conducted.

The state of the model at the beginning of period t, denoted by S_t, consists of three components, i.e. (H_t, V_t, b_t).

One-period cost

The cost incurred in period t, denoted by C_t, consists of three components: measure cost (implementation and maintenance costs), inspection cost, and audit outcome cost. In a non-auditing period t, when the manager decides to take actions A_t^M and a_t^I, the cost incurred is

$$C_t = \sum_{i=1}^{M} c_i^M \, Ind((A_t^M)_i (V_t)_i = -1)_i + \sum_{i=1}^{M} y^M \, Ind\left((V_t)_i = 0 \text{ and } (A_t^M)_i = 1\right) + \sum_{j=1}^{I} c_j^I \, Ind(a_t^I = j),$$

where *Ind(x)* is an indicator function that yields value 1 if condition x is true and 0 otherwise. The first term in the above equation represents the aggregate implementation cost. The second term represents the aggregate maintenance cost while the third term represents the inspection cost.

In an auditing period t, the audit outcome cost incurred is random and depends on the auditing result. The expected cost in period t is

$$E[C_t] = \sum_{i=1}^{M} c_i^M \, Ind((A_t^M)_i (V_t)_i = -1)_i + \sum_{i=1}^{M} y^M \, Ind\left((V_t)_i = 0 \text{ and } (A_t^M)_i = 1\right)$$
$$+ \sum_{j=1}^{I} c_j^I \, Ind(a_t^I = j) + \left(g(H_t)(1-b_t)\right) E_{F^U}\left[z\left(U(b_t), H_t\right)\right]$$

The forth term in the above equation represents the expected audit outcome cost, i.e., the cost induced by the given compliance level in *t*.

Recursion

We define a cost-to-go function $L_t(H_t, V_t, b_t)$ as the expected present value of the cost from period t to the end of the horizon T, when the manager optimally manages the compliance risk, and when the current state at the beginning of period t is (H_t, V_t, b_t). We denote the one-period discount factor as γ. The dynamic programming recursion can be written as follows. In a non-auditing period t, the cost-to-go function is

$$L_t(H_t, V_t, b_t) = \min_{A_t^M, a_t^I} \left\{ \begin{array}{l} \sum_{i=1}^{M} c_i^M \, Ind((A_t^M)_i (V_t)_i = -1)_i + \sum_{i=1}^{M} y^M \, Ind\left((V_t)_i = 0 \text{ and } (A_t^M)_i = 1\right) \\ \quad + \sum_{j=1}^{I} c_j^I \, Ind(a_t^I = j) \\ \quad + \gamma E_{F_i^O}\left[L_{t+1}\left(H_t, V_{t+1}, f(V_t, V_{t+1}, a_t^I, b_t)\right)\right] \end{array} \right\}$$

where

$$\left(V_{t+1}\right)_i = \left\{ \begin{array}{ccc} r_i^M & \text{if} & (A_t^M)_i (V_t)_i = -1 \\ \max\{0, (V_t)_i - 1\} & \text{if} & \left(A_t^M\right)_i = 1 \text{ and } (V_t)_i \neq -1 \\ -1 & \text{if} & \left(A_t^M\right)_i = 0 \end{array} \right.$$

In an auditing period t, the cost-to-go function becomes

$$L_t(H_t, V_t, b_t) = \min_{A_t^M, a_t^I} \left\{ \begin{array}{l} \sum_{i=1}^{M} c_i^M \, Ind((A_t^M)_i (V_t)_i = -1)_i + \sum_{i=1}^{M} y^M \, Ind\left((V_t)_i = 0 \text{ and } (A_t^M)_i = 1\right) \\ \quad + \sum_{j=1}^{I} c_j^I \, Ind(a_t^I = j) \\ \quad + \gamma\left(1 - \left(g(H_t)(1-b_t)\right)\right) \\ \qquad E_{F_i^O}\left[L_{t+1}\left(H_t \cup 1, V_{t+1}, f(V_t, V_{t+1}, a_t^I, b_t)\right)\right] \\ \quad + \gamma\left(g(H_t)(1-b_t)\right) \\ \qquad E_{F^U, F_i^O}\left[z\left(U(b_t), H_t\right) + L_{t+1}\left(H_t \cup U(b_t), V_{t+1}, f(V_t, V_{t+1}, a_t^I, b_t)\right)\right] \end{array} \right\}$$

where

$$H_t \cup x \quad = \quad \{x, (H_t)_1, ..., (H_t)_{K-1}\}$$

The boundary condition of the program is $L_{\tau+1}(H_t, V_t, b_t) = 0$.

Solution

We solve the dynamic program using backward induction algorithm implemented in Java.

In the following section, we give a case study and show the optimal compliance management policy obtained from our dynamic programming model.

3. Case Study

We now provide a simple case study to explain how the introduced model can be used. We will first introduce a number of assumptions. We then present our results.

Assumptions

For the purpose or our case study, we introduce *JustStarted, Inc.*, a medium-sized credit card provider with the following characteristic data:

Table 1 - Information on *JustStarted, Inc.*

Name	JustStarted Inc.
Location	Switzerland
Company size	100
# Customer account managers	10
# Transaction handling managers	20
Customer base	100000

We assume that *JustStarted, Inc.* is affected by a new privacy regulation that includes the following two requirements:

1) *Implement role-based access control to protect and ensure the integrity of electronic data and thus respect customers' privacy.*
2) *Implement mechanisms that ensure that customer data have good quality and are up-to-date.*

To address the above requirements, we further suppose the availability of the following types of measures. Each possible measure is associated with implementation and maintenance costs, a maximally achievable compliance level, and a certain implementation time (cf. Table 2). The figures in Table 2 are estimates based on experience and the contextual data assumed about *JustStarted, Inc.*, which are given in Table 1:

Table 2 - Available compliance measures

	Symbol				aa	ab	ab
Measure#	Measures	Implementation cost	Monthly maintenance cost	Implementation period (months)	% Effectiveness on requirement 1, on the first day if only one implemented	% Effectiveness on requirement 2, on the first day if only one implemented	% Effectiveness on requirement 2 with measure 16
1	6-letter password for every individual user, 3-month forced change	35,000	4,000	0	50%	10%	10%
2	Fingerprint reader access	150,000	700	1	99%	20%	20%
3	Manual plausibility checks/ review of data	96,000	34,000	1	0%	65%	89.50%
4	Update data per customer mail request (letter with signature)	843,333	20,833.00	1	0%	80%	94.00%
5	Address change verification letter (sending to old address)	15,000	2,400.00	0	0%	N/A	0%

Table 2 contains two main classes of measures. Measures 1 and 2 primarily address requirement 1 and measures 3 and 4 primarily address requirement 2. Measure 5 is of a special type in that it does not have a direct effect on any of the two requirements if it is implemented alone. However, if it is implemented together with measures of the second

class, the combined effectiveness on requirement 2 is increased. The new effect is shown in the last column.

Depending on the concrete measure selected, implementation costs may include:

- IT implementation cost, initial user training of the system,
- cost for the preparation of handbooks or guidelines,
- customer training for using the system, and
- loss of customers due to more cumbersome system use.

Likewise, maintenance costs may include costs for:

- time spent for customer service,
- time for fixing bugs of the IT system,
- continuous usage training of the system, and
- administration of user passwords.

In order to monitor compliance and to evaluate the effectiveness of the implemented measures, *JustStarted Inc.* may perform internal inspections. There are three inspection types, ranging from sampling with a low coverage to full inspection covering all implemented measures. The inspection types together with associated costs and related improvements are summarized in Table 3. While the improvement factor of the inspection type 1 is assumed to be certain at 25%, the improvement factors of the other inspection types are assumed to be random. For example, inspection type 2 has two possible improvement factors, which are 70% with a 0.5 probability and 75% with a 0.5 probability. It is also possible not to inspect at all, generating zero inspection cost and yielding no improvement.

Table 3 - Inspection types

inspec#	inspection type	cost per inspection	improvement
-1	No inspection	0	0%
0	Sampling with 5% coverage	400	25%
1	Sampling with 50% coverage	3500	70% w.p 0.5, 75% w.p. 0.5
2	Full inspection (100% coverage)	6000	95% w.p. 0.4, 100% w.p. 0.6

We assume that every *four* periods, *JustStarted Inc.* is being audited. There are four classes of outcomes, which roughly represent the categories used to evaluate operational effectiveness for Sarbanes-Oxley (13): *Full compliance testified, minor control deficiency detected, significant deficiency detected, material weakness found*.

The meaning of these outcomes is as follows:

0. *Full compliance testified*. Auditors have testified that enterprise is fully compliant with all relevant regulatory requirements. The implemented measures address and fulfill all requirements to the fullest satisfaction of the auditors.
1. *Minor deficiency detected*. Auditors have identified minor deficiencies in how requirements have been implemented or with respect to the effectiveness of the implemented measures. This can indicate that a necessary measure is missing, an

existing measure is not properly designed, or a properly implemented measure does not operate as designed.

2. *Significant deficiency detected*. Auditors have detected a significant deficiency, which can be a minor deficiency in a significant measure or an aggregation of such deficiencies that could result in a violation of a relevant requirement that is more than inconsequential.
3. *Material weakness found*. A material weakness is a significant deficiency or an aggregation of significant deficiencies that preclude the implemented measures from providing reasonable assurance that compliance with regulatory requirements can be achieved. The inability to provide such reasonable assurance results from one or more significant deficiencies. The existence of a material weakness precludes the responsible party from concluding that the implemented measures are effective.

The outcome of the audit depends on the compliance level at the time of the audit. As auditing includes some uncontrollable degree of subjectivity (mood of the auditor, context of audit, state of nature, etc.), except for the two degenerate cases of complete compliance and complete non-compliance, the actual compliance level maps to an audit outcome only with a certain probability. For our current example, we assume the following mappings from compliance levels to outcomes, where prob(*i*) denotes the probability with which outcome *i* will be realized (*i={0,1,2,3}*).

Table 4 - Likelihood of audit outcomes

compliance level	prob(0)	prob(1)	prob(2)	prob(3)
100%	1	0	0	0
90% - 99%	0.9	0.1	0	0
70% - 89%	0.1	0.5	0.3	0.1
50% - 69%	0.05	0.25	0.5	0.2
30% - 49%	0	0.1	0.4	0.5
1% - 29%	0	0	0.1	0.9
0%	0	0	0	1

To calculate the audit outcome cost for low compliance with a given regulatory requirement as a function of the current audit outcome and the historical audit outcomes, we assume that *k=2*. Hence only the current outcomes plus the two previous audit outcomes are considered when calculating the current cost resulting from the current audit. Formally,

$$d_t = z(h_0, h_1, h_2) = \left(f_1(h_0) + f_3(h_0) f_2\left(\frac{(h_1 + h_2)}{2} \right) \right) \varphi$$

The penalty factor φ is a constant, which we set to be 1,000,000 in our example. The functions f_1, f_2, and f_3 are defined for all possible values of h_0, h_1, h_2 and $(h_1 + h_2)/2$ as follows:

Table 5 - Calculating audit outcome costs

h_0	$f_1(h_0)$	$f_3(h_0)$
1	0	0
2	10	1
3	30	1.5
4	60	2

φ	1,000,000.00

$(h_1 + h_2)/2$	$f_2((h_1 + h_2)/2)$
1	0
1.5	3
2	6
2.5	10
3	13
3.5	16
4	20

In the same way that the audit outcome cost of a specific regulatory requirement depends on the historical audit outcomes, also the scrutiny with which auditors inspect the compliance status of an enterprise depends on the historical audit outcomes. In Table 6, we present the auditor coverage $q_t = g\left((2h_1 + h_2)/2\right)$ as a function of the previous audit outcomes h_1 and h_2:

Table 6 - Auditor scrutiny as a function of historical audit outcomes

$(2*h_1 + h_2)/3$	$q((2*h_1 + h_2)/3)$
1.00	5.00%
1.33	10.00%
1.67	15.00%
2.00	20.00%
1.33	25.00%
1.67	30.00%
3.00	35.00%
1.33	40.00%
1.67	45.00%
4.00	50.00%

Results

We have solved the above problem using our dynamic and risk-based compliance program using a Java implementation of the presented algorithm. Calculating the optimal solution required approximately 2 hours on a Windows machine with an Intel Pentium 4 CPU 3.00GHz and 3GB of RAM. Assuming a time horizon τ of 60 periods, with audits every 4 periods, and a decay factor ρ of 0.98, the program resulted in four database tables with 93,552 records each (one for each inter-audit period), which we evaluated using SQL queries.

Under the assumptions that the first audit is conducted in the fourth period (t=3), and that no measure has been implemented so far, the program calculated the optimal portfolio of measures (a^{m1} through a^{m5}) that need to be implemented in the first period. It also determined the optimal inspection type a^i for the given setting. The result is shown in Table 7.

Table 7 - Portfolio of optimal measures for the first period (t=0)

h_1	h_2	v_1	v_2	v_3	v_4	v_5	b_t	c_t	a^{m1}_t	a^{m2}_t	a^{m3}_t	a^{m4}_t	a^{m5}_t	a^{i}_t
0	0	-1	-1	-1	-1	-1	0	1475101.6	0	0	0	0	0	-1

Table 7 informs the management of *JustStarted, Inc.* that the optimal portfolio of measures in the starting period includes no measures and thus does not require any inspection yet (a^{i}_t=-1). Following this recommendation and not implementing any measure yet, *JustStarted, Inc.* will find itself in the situation depicted in Table 8 in the next period (t=1). Now, Table 8 informs the responsible management that measures 2 and 3 ($a^{m2}_t = a^{m3}_t$=1) must be implemented and that still no inspection is required (a^{i}_t=-1).

Table 8 - Portfolio of optimal measures in second period (t=1)

h_1	h_2	v_1	v_2	v_3	v_4	v_5	b_t	c_t	a^{m1}_t	a^{m2}_t	a^{m3}_t	a^{m4}_t	a^{m5}_t	a^{i}_t
0	0	-1	-1	-1	-1	-1	0	1490001.6	0	1	1	0	0	-1

In the next period (t=2), the implementation of measures 2 and 3 has not yet finished as both have an implementation period of 1. Table 9 now requires the implementation of an additional measure, namely measure 5, while measures 2 and 3 are maintained. In addition, a full inspection is mandated (a^{i}_t=2). As the new measure has an implementation period of 0, all three measures will be effective in the next period and simultaneously affect *JustStarted, Inc.*'s compliance level.

Table 9 - Portfolio of optimal measures for third period (t=2)

h_1	h_2	v_1	v_2	v_3	v_4	v_5	b_t	c_t	a^{m1}_t	a^{m2}_t	a^{m3}_t	a^{m4}_t	a^{m5}_t	a^{i}_t
0	0	-1	1	1	-1	-1	0	1256567.2	0	1	1	0	1	2

The next period (t=3) is an auditing period. All measures that have already been implemented will affect the compliance level and represent the basis for the audit. The portfolio of optimal measures that are currently implemented now includes measures 2, 3, and 5, resulting in a compliance level of 0.94. The full inspection makes sure that the combined effect of the implemented measures on the compliance level equals the target compliance level of the respective measures.

Table 10 - Portfolio of optimal measures for fourth period (t=3, audit)

h_1	h_2	v_1	v_2	v_3	v_4	v_5	b_t	c_t	a^{m1}_t	a^{m2}_t	a^{m3}_t	a^{m4}_t	a^{m5}_t	a^{i}_t
0	0	-1	0	0	-1	0	0.94	1311975.5	0	1	0	0	1	-1

According to Table 4, with an initial compliance level between 90% and 99%, auditors will attest full compliance (audit outcome 0) with a probability of 0.9 and they will detect minor deficiencies (audit outcome 1) with a probability of 0.1. Hence, from now on *JustStarted, Inc.* might be in either of the two states. As suggested by Table 10, for the next period, only measures 2 and measures 5 need to be maintained and no inspection is conducted. Applying the decay factor of 0.98 to the target compliance level attained through the implemented measures, *JustStarted, Inc.* ends up with a compliance level of roughly 0.58. This is depicted in Table 11. Depending on the audit outcome, h_1 is either 0

or 1. While ex ante we cannot be certain which state is realized, the portfolio of optimal measures for the subsequent period stays the same. *JustStarted, Inc.* should maintain measures 2 and 5 and does not need to conduct an internal inspection in both states.

Table 11 - Portfolio of optimal measures in first period after audit (t=4)

h_1	h_2	v_1	v_2	v_3	v_4	v_5	b_t	c_t	a^{m1}_t	a^{m2}_t	a^{m3}_t	a^{m4}_t	a^{m5}_t	a^{i}_t
0	0	-1	0	-1	-1	0	0.58	1318696	0	1	0	0	1	-1
1	0	-1	0	-1	-1	0	0.58	1334034	0	1	0	0	1	-1

Assuming a previous audit outcome of 0 and maintaining measures 2 and 5 without inspection as suggested by Table 11, in the second period after the audit, *JustStarted, Inc.* reaches the situation summarized in Table 12. With the suggestion to re-implement measure 3, while maintaining measures 2 and 5 without inspection, it is now easy to see that *JustStarted, Inc.* now reaches a compliance state that oscillates between 0.58 and 0.94. This is depicted in Figure 1. Whenever there is an audit period, the three measures 2, 3, and 5 will be in effect, a full inspection will ensure that the compliance level equals the target compliance level and there is a high likelihood (i.e., 0.9) of auditors testifying full compliance.

Table 12 - Portfolio of optimal measures in second period after audit (t=5), assuming h_1=0

h_1	h_2	v_1	v_2	v_3	v_4	v_5	b_t	c_t	a^{m1}_t	a^{m2}_t	a^{m3}_t	a^{m4}_t	a^{m5}_t	a^{i}_t
0	0	-1	0	-1	-1	0	0.57	1328884.9	0	1	1	0	1	-1

However, assuming a previous audit outcome h_1 of 1 and maintaining measures 2 and 5 without inspection, *JustStarted, Inc.*'s decisions in period 5 are summarized in Table 13.

Table 13 - Portfolio of optimal measures in second period after audit (t=5), assuming h_1=1

h_1	h_2	v_1	v_2	v_3	v_4	v_5	b_t	c_t	a^{m1}_t	a^{m2}_t	a^{m3}_t	a^{m4}_t	a^{m5}_t	a^{i}_t
1	0	-1	0	-1	-1	0	0.57	1344377.8	0	1	1	0	1	-1

Assuming that the audit outcome was 1, in the next period, *JustStarted, Inc.* still maintains measures 2 and 5, is implementing measure 3, does not inspect and thus finds itself in the situation summarized in Table 14.

Table 14 - Portfolio of optimal measures in period t=6, assuming h_1=1

h_1	h_2	v_1	v_2	v_3	v_4	v_5	b_t	c_t	a^{m1}_t	a^{m2}_t	a^{m3}_t	a^{m4}_t	a^{m5}_t	a^{i}_t
1	0	-1	0	1	-1	0	0.56	1255356.2	0	1	1	0	1	1

The next period (t=7) is again an auditing period. The re-implementation of measure 3 has been completed and *JustStarted, Inc.* again reaches a compliance level of 0.94. We summarize the respective information in Table 15. We may now recognize an interesting result: Also if we assume the worse audit outcome (i.e., h_1=1), it is optimal to re-implement measure 3, while maintaining measures 2 and 5, and thereby achieve a compliance level of 0.94 again. As long as *JustStarted, Inc.* manages to attain this

compliance level in the auditing periods, it will never experience an audit outcome lower than 1. Given the available measures and other assumptions of this case study, the worst possible audit outcome is that the auditors register minor deficiencies and that *JustStarted, Inc.* needs to implement their recommendations. In such cases, and in general with a track of subsequent audit outcomes of 1, the cost of compliance is slightly higher than in the case where *JustStarted, Inc.* reaches the audit outcome 0 (which is much more likely in any case).

Table 15 - Portfolio of optimal measures in t=7 (audit), assuming h_1=1

h_1	h_2	v_1	v_2	v_3	v_4	v_5	b_t	c_t	a^{m1}_t	a^{m2}_t	a^{m3}_t	a^{m4}_t	a^{m5}_t	a^{i}_t
1	0	-1	0	0	-1	0	0.94	1326249	0	1	0	0	1	-1

By thus reconstructing the evolution of the compliance level and cost of compliance of the fictitious company *JustStarted, Inc.*, we are able to understand that attaining a high level of compliance with regulatory requirements may not only be a moral obligation but may also be economically optimal.

As an additional result, Figure 2 shows how the optimal inspection type for any given compliance level varies with the historical audit outcomes in a period before just before an audit. With the current target compliance level being 0.58 and the previous audit outcome being 1, *JustStarted, Inc.* would almost always conduct a full inspection (i.e., a^{i}_t=2) to raise its compliance level up to the target level. Only in cases where the current compliance level is already close to the target compliance level, a partial inspection with 50% coverage (i.e., a^{i}_t=1) will be enough. In case the previous audit outcome was 0, *JustStarted, Inc.* does not require an equally pronounce inspection strategy. For example, given a compliance level of 0.5, *JustStarted, Inc.* in this would conduct an inspection with only 50% coverage, as opposed to full inspection if it had an audit outcome of 1.

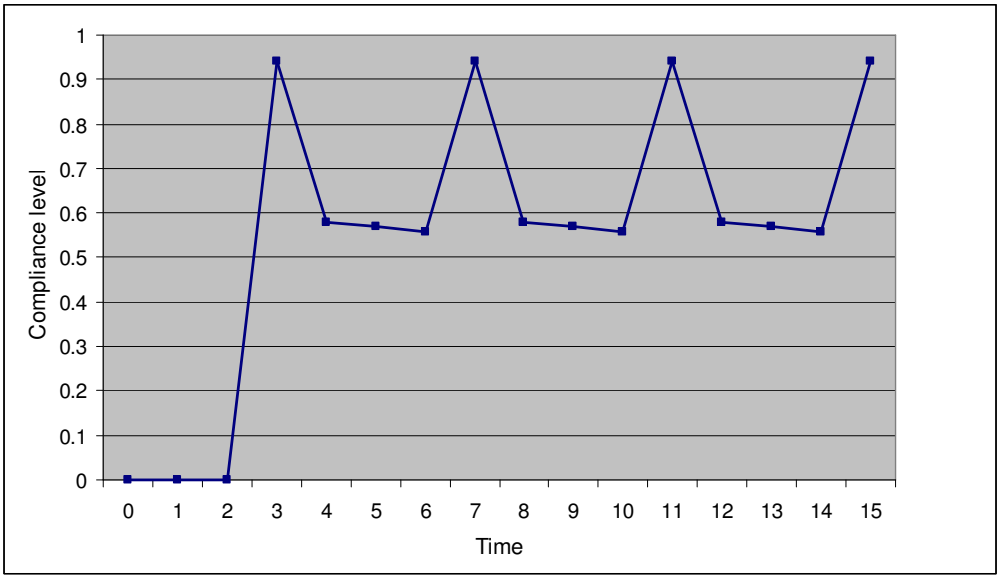

Figure 1 - Evolution of compliance level over time.

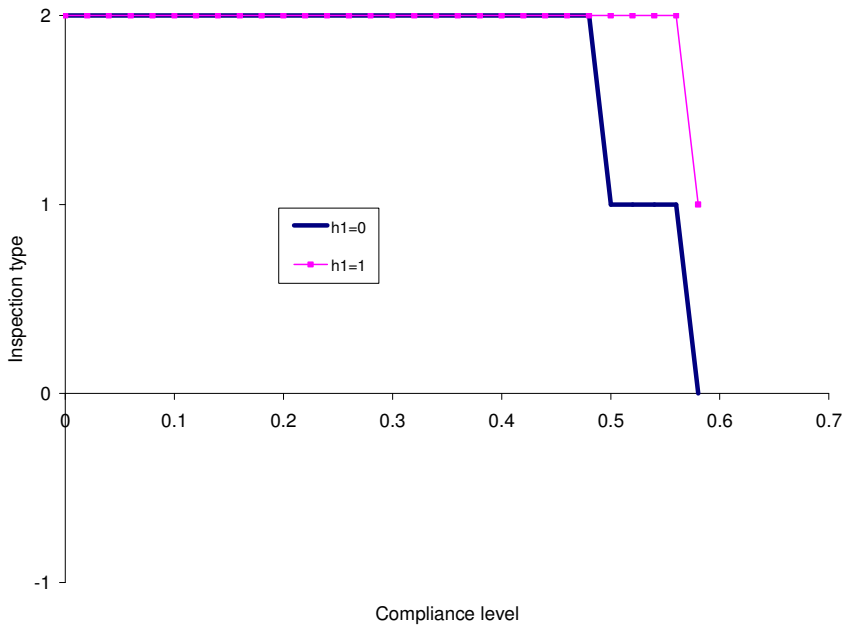

Figure 2 - Optimal inspection type depending on compliance level in two possible audit outcome states (i.e., h1=0 and h1=1 for a given Vt=[-1, 0, 1, -1, 0]).

4. Conclusion

For large enterprises, attempting to reach perfect compliance with abundant regulatory requirements may be idealistic but close to impossible. Trying to reach this sublime goal potentially consumes more resources than is economically optimal. In this paper we have described a dynamic and risk-based approach to compliance management based on dynamic programming. Given a set of available measures to address compliance concerns, our approach determines the portfolio of optimal measures that need to be implemented and the optimal type and frequency of internal inspection at any point in time within the decision horizon.

Our main contribution lies in the new way of looking at the compliance management problem, specifically in our approach to address and solve it. We want to stress the notion of compliance as a continuous phenomenon rather than regard it as a binary property. Hence, compliance needs to be measured on a ratio scale and is to be managed in a risk-based way to support the prioritization and optimal selection of measures and inspection frequency.

We are aware of the fact that it may be difficult to populate our model with meaningful input data. We thus would like to stress a number of considerations. While it may often not be possible to come up with precise estimates of various input parameters or function specifications grounded on solid empirical data, our tool still lends itself very nicely for sensitivity analysis and scenario-based decision evaluation. It may therefore prove to be a valuable decision support system when contemplating feasible steps in the process of managing enterprise compliance. In this context, a certain level of imprecision when estimating individual model parameters may well be tolerable.
What is more, we are optimistic that today's enterprises will further improve with respect to data integration through standards, harmonization, and simplification. We also observe that more and more IT systems are being instrumented to allow for event monitoring. Over time, we thus expect to see enterprises evolving ever more towards a point where continuous monitoring and assurance are within reach and our quantitative model can be populated with more reliable data. One can conceive of enterprises belonging to the same industry sharing compliance risk-relevant input data (e.g., on audit outcome cost, measure effectiveness, audit coverage, etc.) in an anonymous form, similar to consortiums of financial institutions, such as Operational Riskdata eXchange Association (ORX) (18), sharing operational risk data anonymously.

Being interested in effective regulation, also governmental and rule-making institutions might take advantage of our model. Using our approach and assuming reliable input data, law makers could better evaluate whether a new regulation can be effectively enforced by simulating enterprises' reactions to the new regulation. Governments would thus be in a position to minimize bureaucratic overhead, avoid ineffective regulation and guide optimal behavior with respect to regulatory compliance.

Cited references (and notes)

1. “Sarbanes-Oxley Act of 2002”, PL 107-204, 116 Stat 745 (2002).
2. “USA Patriot Act of 2001”, PL 107-56, HR 3162 RDS (2001).
3. P. L. Bernstein. “Against The Gods – The Remarkable Story of Risk”, John Wiley and Sons, New York (1996).
4. A. J. McNeil, R. Frey, and P. Embrechts, “Quantitative Risk Management: Concepts, Techniques, and Tools”, Princeton University Press (2005).
5. J. C. Hull, “Options, Futures and other Derivative Securities”, 2nd Edition. Prentice-Hall. Englewood Cliffs, New Jersey (1993).
6. P. J. Schönbucher, “Credit Derivatives Pricing Models: Models, Pricing, Implementation”, Wiley Finance (2003)
7. G. E. G. Beroggi and W. A. Wallace, “Operational Risk Management - A New Paradigm for Decision-Making”, IEEE Transactions on Systems, Man and Cybernetics, Vol. 24, No. 10, pp. 1450-1457 (1994).
8. M. Leippold and P. Vanini, “The Quantification of Operational Risk”, Journal of Risk, Vol. 8, No. 1, pp. 59-85 (2005).
9. C. Supatgiat, C. Kenyon, and L. Heusler, “Cause-to-Effect Operational Risk Quantification”; Risk Management: an International Journal (2005).
10. Ashish Gehani and Gershon Kedem, “RheoStat: Real-Time Risk Management”, In Proceedings of Recent Advances in Intrusion Detection: 7th International Symposium, RAID 2004, (2004)
11. Lam-for Kwok and Dennis Longley: Security Modelling for Risk Analysis. SEC 2004, 29-46 (2004)
12. Committee of Sponsoring Organizations of the Treadway Commission (COSO), “Internal Control -- Integrated Framework”, Jersey City, NJ: AICPA/COSO (1992).
13. IT Governance Institute, “Control Objectives for Information and Related Technology (COBIT)”, Version 4.0, (2005).
14. Office of Government Commerce (OGC), “IT Infrastructure Library (ITIL)” (2006)
15. IT Governance Institute, “IT Control Objectives for Sarbanes-Oxley” (2004).
16. R. Bellman, *Dynamic Programming*. Princeton Univ. Press, Princeton, NJ (1957)
17. E. Denardo, *Dynamic Programming : Models and Applications*, Dover Publications (2003)
18. Operational Riskdata eXchange Association (ORX), http://www.orx.org/

Biographies

Samuel Müller, IBM Research Division, IBM Zurich Research Laboratory, Säumerstrasse 4, 8803 Rüschlikon, Switzerland (sml@zurich.ibm.com). Mr. Müller obtained a M.S. in Computer Science from the University of Zurich in 2004 and a M.A. in Economics from the University of Zurich in 2006. He joined IBM Research in Zurich in 2004, where he is currently doing research in the area of risk and compliance. In parallel, he is working towards his doctorate as an external Ph.D. student at the Swiss Federal Institute of Technology (ETH) Zurich, where he is a member of the Information Security group. His thesis advisors are Prof. Dr. David Basin and Prof. Dr. Birgit Pfitzmann and his research interests include modal logics, formal methods, modeling methodology, risk & compliance management, game theory and economics.

Chonawee Supatgiat, IBM Research Division, IBM Zurich Research Laboratory, Säumerstrasse 4, 8803 Rüschlikon, Switzerland (csu@zurich.ibm.com). Dr. Supatgiat is a research staff member in Business Optimization group at IBM Zurich Research Laboratory, Switzerland. Chonawee received a B.S.E. degree in Industrial Engineering from Chulalongkorn University, Thailand, in 1993, and M.S.E. and Ph.D. degrees in Industrial and Operations Engineering from the University of Michigan, Ann Arbor, in 1996 and 1999, respectively. His research interests include sequential decision-making processes, large-scale stochastic optimization, financial engineering, and game theory.