


Editors

RON S. KENETT

YOSSI RAANAN

Operational Risk Management

A Practical Approach to
Intelligent Data Analysis



 WILEY

STATISTICS IN PRACTICE

Foreword

The recognition from the Basel Committee of Banking Supervisors of operational risks as a separate risk management discipline has promoted in the past years intense and fruitful discussions, both inside and outside the banking and financial sectors, on how operational risks can be managed, assessed and prevented, or at least mitigated.

However, for several reasons, including the fact that operational risks appear at the same time multifaceted and of a somewhat indefinite shape, inadequate attention has been given so far to what operational risks really are, and to how they can be correctly identified and captured.

Indeed, the first objective of a risk management programme is to identify clearly the playing field to where investments and resources should be directed. This is even more important for operational risk management, since its scope crosses all industry sectors and all types of firms and the fact that it essentially originates from those variables that constitute the heart of any organization: people, processes and systems.

This book attempts to give an appropriate level of attention to this significant topic by using an interdisciplinary, integrated and innovative approach.

The methodologies and techniques outlined here, reading ‘behind and beyond’ operational risks, aim to move forward in the interpretation of this type of risk and of the different ways it can reveal. The objective of capturing knowledge on operational risk, rather than just information, is crucial for the building of sound processes for its management, assessment and prevention or mitigation.

Another noteworthy feature of this work is the effort – pursued by providing practical examples of implementation of an operational risk framework (or part of it) in different industry sectors – to demonstrate how concepts, techniques and methodologies developed in a specific field, for the handling of operational risks, can be adopted in (or adapted to) other industrial domains. If considered all together, these aspects can significantly contribute to make this discipline evolve towards high, sustainable and convergent standards and, above all, to change its nature from (a bit less) ‘art’ to (a bit more) ‘science’, which, in the end, is the ultimate objective that all operational risk managers are trying to achieve.

Marco Moscadelli
Bank of Italy and Committee of European Banking Supervisors

Preface

This book is a result of the MUSING (Multi-industry, Semantic-based, next generation business INtelligence) project collaborative effort, an R&D venture co-funded by the European Commission under the FP6 Information Society Technology Programme. The project covered a period of four years, witnessing many dramatic events, including the Wall Street crash in September 2008. It was designed to be driven by customer requirements in three main areas of application: operational risk management, financial risk management and internationalization. The idea was to develop innovative solutions to customer requirements in these three domains, with partners that are leaders in their fields. The MUSING partners combined expertise and experience in risk management, information extraction and natural language processing, with ontology engineering, data mining and statistical modelling. The focus in this book is operational risk management. The customers we had in mind are financial institutions implementing Basel II regulations, industrial companies developing, manufacturing and delivering products and services, health care organizations and others with exposure to operational risks with potential harmful effects and economic impact.

The key visionaries behind the project were Roberto Gagliardi and Paolo Lombardi. At the inaugural meeting, on 5 April 2006 in Pisa, Italy, they presented a slide with the project components organized in the shape of a sailboat. The basic research and integration partners were responsible for the keel and cockpit. The sails, pushing the boat forward, were the three application areas. This strategy, where customers and users motivate researchers to develop innovative solutions based on state-of-the-art technologies, is what MUSING was about.

Unfortunately Roberto passed away as the boat started sailing, but the MUSING vision kept the project on track.

Operational risk management is a complicated topic. Its precise definition is elusive and its ‘boundary of meaning’ has evolved over time. From a classification of residual risks, that is risks not identified as financial risks or market risks, it has become a key area with specific methods, dedicated technology and dedicated indicators and scoring systems. This book covers many of the state-of-the-art techniques in this area, with many implementation examples from real life from MUSING. Some chapters are more mathematical than others, some are fully descriptive. In designing the book we wanted to balance the various disciplines involved in setting up an infrastructure for modern operational risk management.

The creative synergy between experts in various disciplines has made MUSING a unique project. We hope the book will convey this message to the readers. Not all the authors who contributed to the book were part of the MUSING project. All chapters, however, present the latest advances in operational risk management by a combination of novel methods in the context of real problems, as envisaged in the MUSING project. As such, we believe that the book provides a solid foundation and challenging directions for operational risk management.

In preparing an edited volume, it is natural for many people to be involved. As editors, it was a pleasure and a privilege to work with the authors of the 14 chapters. These authors were also kind enough to serve as internal reviewers. The typos and mistakes that sneaked in remain, however, our responsibility. We want to thank the authors who dedicated their time and talent to write these chapters and all our colleagues in the MUSING project who helped develop this knowledge. Special thanks are due to the project coordinators Marcus Spies and Thierry Declerk, the Metaware team, the project reviewers, Professors Vadim Ermolayev, Mark Lycett, Aljosa Pasic and the project officer, Francesco Barbato – they all contributed significantly to the success of MUSING. The help of Drs Emil Bashkansky and Paolo Lombardi in reviewing the chapters is also gratefully acknowledged. Finally, we would like to thank Dr Ilaria Meliconi, Heather Kay and Richard Davies from John Wiley & Sons, Ltd for their help, directions and patience.

Ron S. Kenett and Yossi Raanan

Introduction

Operational risk management is becoming a key competency for organizations in all industries. Financial institutions, regulated by the Basel II Accord, need to address it systematically since their level of implementation affects their capital requirements, one of their major operational expenses. Health organizations have been tackling this challenge for many years. The Institute of Medicine reported in 2000 that 44 000–98 000 patients die each year in the United States as a result of medication errors, surgical errors and missed diagnoses, at an estimated cost to the US economy of \$17–29 billion. Operational risks affect large organizations as well as small and medium-sized enterprises (SMEs) in virtually all industries, from the oil and gas industry to hospitals, from education to public services.

This multi-author book is about tracking and managing operational risks using state-of-the-art technology that combines the analysis of qualitative, semantic, unstructured data with quantitative data. The examples used are mostly from information technology but the approach is general. As such, the book provides knowledge and methods that can have a substantial impact on the economy and quality of life.

The book has four main parts. Part I is an introduction to operational risk management, Part II deals with data for operational risk management and its handling, Part III covers operational risk analytics and Part IV concludes the book with several applications and a discussion on how operational risk management integrates with other disciplines. The 14 chapters and layout of the book are listed below with short descriptions.

Part I: Introduction to Operational Risk Management

This first part of the book is introductory with a review of modern risk management in general and a presentation of specific aspects of operational risk management issues.

Chapter 1: *Risk management: a general view* (R. Kenett, R. Pike and Y. Raanan)

This chapter introduces the concepts of risk management and positions operational risk management within the overall risk management landscape. The topics covered include definitions of risks, aspects of information quality and a discussion of state-of-the-art enterprise risk management. The organizations the

authors have in mind are financial institutions implementing Basel II regulations, industrial companies developing, manufacturing and delivering products and services, health care services and others with exposure to risks with potential harmful effects. The chapter is meant to be a general introduction to risk management and a context-setting background for the 13 other chapters of the book.

Chapter 2: *Operational risk management: an overview* (Y. Raanan, R. Kenett and R. Pike)

This chapter introduces the general concepts of operational risk management in the context of the overall risk management landscape. Section 2.2 provides a definition of operational risk management, Section 2.3 covers the key techniques of this important topic, Section 2.4 discusses statistical models and Section 2.5 covers several measurement techniques for assessing operational risks. The final section summarizes the chapter and provides a roadmap for the book.

Part II: Data for Operational Risk Management and its Handling

Operational risk management relies on diverse data sources, and the handling and management of this data requires novel approaches, methods and implementations. This part is devoted to these concepts and their practical applications. The applications are based on case studies that provide practical, real examples for the practitioners of operational risk management.

Chapter 3: *Ontology-based modelling and reasoning in operational risks* (C. Leibold, H.-U. Krieger and M. Spies)

This chapter discusses the design principles of operational risk ontologies for handling semantic unstructured data in operational risk management (OpR). In particular, the chapter highlights the contribution of ontology modelling to different levels of abstraction in OpR. Realistic examples from the MUSING project and application-domain-specific ontologies are provided. A picture is drawn of axiomatic guidelines that provide a foundation for the ontological framework and refers to relevant reporting and compliance standards and generally agreed best practices.

Chapter 4: *Semantic analysis of textual input* (H. Saggion, T. Declerck and K. Bontcheva)

Information extraction is the process of extracting from text specific facts in a given target domain. The chapter gives an overview of the field covering components involved in the development and evaluation of an information extraction system such as parts of speech tagging or named entity recognition. The chapter introduces available tools such as the GATE system and illustrates rule-based approaches to information extraction. An illustration of information extraction in the context of the MUSING project is presented.

Chapter 5: *A case study of ETL for operational risks* (V. Grossi and A. Romei)

Integrating both internal and external input sources, filtering them according to rules and finally merging the relevant data are all critical aspects of business analysis and risk assessment. This is especially critical when internal loss data is not sufficient for effective calculation of risk indicators. The class of tools responsible for these tasks is known as *extract, transform and load (ETL)*. The chapter reviews state-of-the-art techniques in ETL and describes an application of a typical ETL process in the analysis of causes of operational risk failures. In particular, it presents a case study in information technology operational risks in the context of a telecommunication network, highlighting the data sources, the problems encountered during the data merging and finally the solution proposed and implemented by means of ETL tools.

Chapter 6: *Risk-based testing of web services* (X. Bai and R. Kenett)

A fundamental strategy for mitigating operational risks in web services and software systems in general is testing. Exhaustive testing of web services is usually impossible due to unavailable source code, diversified user requirements and the large number of possible service combinations delivered by the open platform. The chapter presents a risk-based approach for selecting and prioritizing test cases to test service-based systems. The problem addressed is in the context of semantic web services. Such services introduce semantics to service integration and interoperation using ontology models and specifications like OWL-S. They are considered to be the future in World Wide Web evolution. However, due to typically complex ontology relationships, semantic errors are more difficult to detect, compared with syntactic errors. The models described in the chapter analyse semantics from various perspectives such as ontology dependency, ontology usage and service workflow, in order to identify factors that contribute to risks in the delivery of these services. Risks are analysed from two aspects, namely *failure probability and importance*, and three layers: ontology data, specific services and composite services. With this approach, test cases are associated to the semantic features and schedule test execution on the basis of risks of their target features. Risk assessment is then used to control the process of web services progressive group testing, including test case ranking, test case selection and service ruling out. The chapter presents key techniques used to enable an effective adaptation mechanism: adaptive measurement and adaptation rules. As a statistical testing technique, the approach aims to detect, as early as possible, the problems with highest impact on the users. A number of examples are used to illustrate the approach.

Part III: Operational Risk Analytics

The data described in Part II requires specialized analytics in order to become information and in order for that information to be turned, in a subsequent phase of its analysis, into knowledge. These analytics will be described here.

Chapter 7: *Scoring models for operational risks* (P. Giudici)

This chapter deals with the problem of analysing and integrating qualitative and quantitative data. In particular it shows how, on the basis of the experience and opinions of internal company 'experts', a scorecard is derived, producing a ranking of different risks and a prioritized list of improvement areas and related controls. Scorecard models represent a first step in risk analysis. The chapter presents advanced approaches and statistical models for implementing such models.

Chapter 8: *Bayesian merging and calibration for operational risks* (S. Figini)

According to the Basel II Accord, banks are allowed to use the advanced measurement approach (AMA) option for the computation of their capital charge covering operational risks. Among these methods, the loss distribution approach (LDA) is the most sophisticated one. It is highly risk sensitive as long as internal data is used in the calibration process. Given that, LDA is more closely related to the actual risks of each bank. However, it is now widely recognized that calibration on internal data only is not sufficient for computing accurate capital requirements. In other words, internal data should be supplemented with external data. The goal of the chapter is to provide a rigorous statistical method for combining internal and external data and to ensure that merging both databases results in unbiased estimates of the severity distribution.

Chapter 9: *Measures of association applied to operational risks* (R. Kenett and S. Salini)

Association rules are basic analysis tools for unstructured data such as accident reports, call-centre recordings and customer relationship management (CRM) logs. Such tools are commonly used in basket analysis of shopping carts for identifying patterns in consumer behaviour. The chapter shows how association rules are used to analyse unstructured operational risk data in order to provide risk assessments and diagnostic insights. It presents a new graphical display of association rules that permits effective clustering of associations with a novel interest measure of association rule called the relative linkage disequilibrium.

Part IV: Operational Risk Applications and Integration with other Disciplines

Operational risk management is not a stand-alone management discipline. This part of the book demonstrates how operational risk management relates to other management issues and intelligent regulatory compliance.

Chapter 10: *Operational risk management beyond AMA: new ways to quantify non-recorded losses* (G. Aprile, A. Pippi and S. Visinoni)

A better understanding of the impact of IT failures on the overall process of operational risk management can be achieved not only by looking at the risk events

with a bottom line effect, but also by drilling down to consider the potential risks in terms of missed business opportunities and/or near losses. Indeed, for banking regulatory purposes, only events which are formally accounted for in the books are considered when computing the operational capital at risk. Yet, the 'hidden' impact of operational risks is of paramount importance under the implementation of the Pillar 2 requirements of Basel II, which expands the scope of the analysis to include reputation and business risk topics. This chapter presents a new methodology in operational risk management that addresses these issues. It helps identify multiple losses, opportunity losses and near misses, and quantifies their potential business impact. The main goals are: (1) to reconstruct multiple-effect losses, which is compliant with Basel II requirements; and (2) to quantify their potential impact due to reputation and business risks (opportunity losses) and low-level events (near misses), which is indeed a possible extension to the Basel II advanced measurement approach (AMA). As a consequence, the proposed methodology has an impact both on daily operations of a bank and at the regulatory level, by returning early warnings on degraded system performance and by enriching the analysis of the risk profile beyond Basel II compliance.

Chapter 11: *Combining operational risks in financial risk assessment scores* (M. Munsch, S. Rohe and M. Jungemann-Dorner)

The chapter's central thesis is that efficient financial risk management must be based on an early warning system monitoring risk indicators. Rating and scoring systems are tools of high value for proactive credit risk management and require solid and carefully planned data management. The chapter introduces a business retail rating system based on the Creditreform solvency index which allows a fast evaluation of a firm's creditworthiness. Furthermore, it evaluates the ability of quantitative financial ratings to predict fraud and prevent crimes like money laundering. This practice-oriented approach identifies connections between typical financing processes, operational risks and risk indicators, in order to point out negative developments and trends, enabling those involved to take remedial action in due time and thereby reverse these trends.

Chapter 12: *Intelligent regulatory compliance* (M. Spies, R. Gubser and M. Schacher)

In view of the increasing needs for regulation of international markets, many regulatory frameworks are being defined and enforced. However, the complexity of the regulation rules, frequent changes and differences in national legislation make it extremely complicated to implement, check or even prove regulatory compliance of company operations or processes in a large number of instances. In this context, the Basel II framework for capital adequacy (soon to evolve to Basel III) is currently being used for defining internal assessment processes in banks and other financial services providers. The chapter shows how recent standards and specifications related to business vocabularies and rules enable intelligent regulatory compliance (IRC). IRC is taken to mean semi-automatic or fully automated

procedures that can check business operations of relevant complexity for compliance against a set of rules that express a regulatory standard. More specifically, the BMM (Business Motivation Model) and SBVR (Semantics of Business Vocabularies and business Rules) specifications by the Object Management Group (OMG) provide a formal basis for representing regulation systems in a sufficiently formal way to enable IRC of business processes. Besides the availability of automatic reasoning systems, IRC also requires semantics-enabled analysis of business service and business performance data such as process execution logs or trace data. The MUSING project contributed several methods of analysis to the emerging field of IRC. The chapter discusses standards and specifications for business governance and IRC based on BMM and SBVR.

Chapter 13: *Democratisation of enterprise risk management* (P. Lombardi, S. Piscuoglio, R. Kenett, Y. Raanan and M. Lankinen)

This chapter highlights the interdisciplinary value of the methodologies and solutions developed for semantically enhanced handling of operational risks. The three domains dealt with are operational risk management, financial risk management and internationalization. These areas are usually treated as ‘worlds apart’ because of the distance of the players involved, from financial institutions to public administrations, to specialized consultancy companies. This proved to be fertile common ground, not only for generating high-value tools and services, but also for a ‘democratised’ approach to risk management, a technology of great importance to SMEs worldwide.

Chapter 14: *Operational risks, quality, accidents and incidents* (R. Kenett and Y. Raanan)

This concluding chapter presents challenges and directions for operational risk management. The first section provides an overview of a possible convergence between risk management and quality management. The second section is based on a mapping of uncertainty behaviour and decision-making processes due to Taleb (2007). This classification puts into perspective so-called ‘black swans’, rare events with significant impact. The third section presents a link between management maturity and the application of quantitative methods in organizations. The fourth section discusses the link between accidents and incidents and the fifth section is a general case study from the oil and gas industry. This illustrates the applicability of operational risk management to a broad range of industries. A final summary section discusses challenges and opportunities in operational risks. Chapter 14 refers throughout to previous chapters in order to provide an integrated view of the material contained in the book.

The book presents state-of-the-art methods and technology and concrete implementation examples. Its main objective is to push forward the operational risk

management envelope in order to improve the handling and prevention of risks. It is hoped that this work will contribute, in some way, to organizations which are motivated to improve their operational risk management practices and methods with modern technology. The potential benefits of such improvements are immense.

Notes on Contributors

Ron S. Kenett

Ron Kenett is Chairman and CEO of KPA Ltd (an international management consulting firm with head office in Raanana, Israel), Research Professor at the University of Turin and International Professor Associate at the Center for Risk Engineering of NYU Poly, New York. He has over 25 years of experience in restructuring and improving the competitive position of organizations by integrating statistical methods, process improvements, supporting technologies and modern quality management systems. For 10 years he served as Director of Statistical Methods for Tadiran Telecommunications Corporation and, previously, as researcher at Bell Laboratories in New Jersey. His 160 publications and seven books are on topics in industrial statistics, multivariate statistical methods, improvements in system and software development and quality management. He is Editor in Chief, with F. Ruggeri and F. Faltin, of the *Encyclopedia of Statistics in Quality and Reliability* (John Wiley & Sons, Inc., 2007) and of the international journal *Quality Technology and Quantitative Management*. He was President of ENBIS, the European Network for Business and Industrial Statistics, and his consulting clients include hp, EDS, SanDisk, National Semiconductors, Cisco, Intel, Teva, Merck Serono, Perrigo, banks, healthcare systems, utility companies and government organizations. His PhD is in mathematics from the Weizmann Institute of Science, Rehovot, Israel, and BSc in mathematics, with first-class honours, from Imperial College, London University.

Yossi Raanan

Yossi Raanan is a Senior Consultant and Strategic Partner in KPA, Ltd and Senior Lecturer at the Business School of the College of Management – Academic Studies in Rishon LeZion, Israel. He is also a former dean of that school. He has extensive experience in areas of information technology, data and computer communications and quality management as well as in applying management concepts, tools and know-how to realistic problems, creating applicable, manageable solutions that improve business performance and profitability. His publications and conference talks mirror this knowledge. In addition, he has served on the

board of directors of a leading mutual trust company and served as the head of its investment committee, as the chairman of the board of a government-owned company, and a director of a number of educational institutions and a number of start-up technology companies. His PhD is in operations research from Cornell University, Ithaca, NY, with a dissertation on ground-breaking applications of the theories of Lloyd S. Shapley and Robert J. Aumann (Nobel laureate in Economics, 2005) to real-life situations. His BSc, *cum laude*, is in mathematics from the Hebrew University, Jerusalem.

Giorgio Aprile

Giorgio Aprile is head of the operational risk management function in the Monte dei Paschi di Siena Banking Group, the third banking group in Italy and the second certified for an AMA model. He started his activities in OpR in 2003, and from 2005 he has been responsible for implementing the Advanced Measuring Approach (AMA) in the MPS group; the AMA model has been fully running since January 2007. He was born in 1971 and graduated in nuclear engineering in 1996 with a dissertation thesis on the safety analysis of Ukrainian nuclear power plants. He worked as a risk analyst in the oil and gas industry for five years, on different research projects, mainly focused on the 'smart' use of field data to reduce operational risks in offshore plants and the prevention of environmental disasters. At the end of 2001 he joined the MPS group and focused on the development of advanced tools for risk management and strategic marketing for the bank. From June 2005 he has been the head of the OpR management function.

Xiaoying Bai

Xiaoying Bai is currently an Associate Professor at the Department of Computer Science and Technology of Tsinghua University. She received her PhD degree in computer science in 2001 from Arizona State University in the United States. After that, she joined the Department of Computer Science and Technology of Tsinghua University in 2002 as an Assistant Professor and was promoted to Associated Professor in 2005. Her major research area is software engineering, especially model-driven testing and test automation techniques in various software paradigms such as distributed computing, service-oriented architecture and embedded systems. She has led more than 10 projects funded by the National Key Science and Technology Project, National Science Foundation and National High Tech 863 Program in China, as well as international collaboration with IBM and Freescale. She was also involved as a key member of the Key Project of Chinese National Programs for Fundamental Research and Development (973 Program). She has published over 50 papers in journals and international conference proceedings of ANSS, COMPSAC, ISADS, SIGSOFT CBSE, ICWS, etc. She is the co-author of a Chinese book, *Service-Oriented Software Engineering*. She was the Program Chair of the first IEEE International Conference on Service Oriented System Engineering and is now serving or has served as PC member for

many software engineering conferences, including ICEBE, QSIC, COMPSAC, SEKE, HASE, WISE, etc., and as a reviewer for international journals.

Kalina Bontcheva

Kalina Bontcheva is a Senior Researcher at the Natural Language Processing Laboratory of the University of Sheffield. She obtained her PhD from the University of Sheffield in 2001 and has been a leading developer of GATE since 1999. She is Principal Investigator for the TAO and MUSING projects where she coordinates works on ontology-based information extraction and ontology learning.

Thierry Declerck

Thierry Declerck (MA in philosophy, Brussels; MA in computer linguistics, Tübingen) is a Senior Consultant at the DFKI Language Technology Lab. Before joining DFKI in 1996, he worked at the Institute of Natural Language Processing (IMS) in Stuttgart. At DFKI, he worked first in the field of information extraction. He was later responsible for the EU FP5 project MUMIS (MultiMedia Indexing and Searching). He has also worked at the University of Saarland, conducting two projects, one on a linguistic infrastructure for e-Content and the other one, Esperanto, on the relation between NLP and the Semantic Web. He has actually led the DFKI contribution to the European Network of Excellence 'K-Space' (Knowledge Space of semantic inference for automatic annotation and retrieval of multimedia content, see www.k-space.eu) and is coordinating the research work packages of the IP MUSING (see www.musing.eu), both projects being part of the 6th Framework Programme in IST. He is also actively involved in standardization activities in the context of ISO TC37/SC4.

Silvia Figini

Silvia Figini has a PhD in statistics from Bocconi University in Milan. She is a researcher in the Department of Statistics and Applied Economics L. Lenti, University of Pavia, a member of the Italian Statistical Society and author of publications in the area of methodological statistics, Bayesian statistics and statistical models for financial risk management. She teaches undergraduate courses in applied statistics and data mining.

Paolo Giudici

Paolo Giudici is Professor at the University of Pavia where he is a lecturer in data analysis, business statistics and data mining, as well as risk management (at Borromeo College). He is also Director of the Data Mining Laboratory; a member of the University Assessment Board; and a coordinator of the Institute of Advanced Studies school on 'Methods for the management of complex systems'. He is the author of 77 publications, among which are two research books and 32

papers in Science Citation Index journals. He has spent several research periods abroad, in particular at the University of Bristol, the University of Cambridge and at the Fields Institute (Toronto) for research in the mathematical sciences. He is the coordinator of two national research grants: one (PRIN, 2005–2006) on ‘Data mining methods for e-business applications’; and one (FIRB 2006–2009) on ‘Data mining methods for small and medium enterprises’. He is also the local coordinator of a European integrated project on ‘Data mining models for advanced business intelligence applications’ (MUSING, 2006–2010) and responsible for the Risk Management Interest Group of the European Network for Business and Industrial Statistics. He is also a member of the Italian Statistical Society, the Italian Association for Financial Risk Management and the Royal Statistical Society.

Valerio Grossi

Valerio Grossi holds a PhD in computer science from the University of Pisa. Currently, he is a Research Associate at the Department of Computer Science, University of Pisa, in the business intelligence research area. He was involved in the MUSING project on the development of strategies of operational risk management. His research activities include machine learning, data mining and knowledge discovery with special reference to mining data streams.

Rolf Gubser

Rolf Gubser graduated in computer science in 1990. He is a founding member of KnowGravity, Inc., a leading contributor to several OMG specifications like BMM (Business Motivation Model), SBVR (Semantics of Business Vocabulary and business Rules) and BPMN (Business Process Modelling Notation). Based on those specifications, he focuses on developing and applying Model Driven Enterprise Engineering™, a holistic approach that integrates strategic planning, risk and compliance management, business engineering, as well as IT support. Before KnowGravity, he worked as a Senior Consultant at NCR and Born Informatik AG.

Monika Jungemann-Dorner

Monika Jungemann-Dorner, born in 1964, has a Masters degree in linguistics and economics. She has been working on international projects for more than 10 years. At the German Chamber Organization she was responsible for the management and deployment of a number of innovative European projects (e.g. knowledge management and marketplaces). Since January 2004, she has worked as Senior International Project Manager for the Verband der Vereine Creditreform eV, Germany.

Hans-Ulrich Krieger

Hans-Ulrich Krieger is a Senior Researcher at the German Research Centre for Artificial Intelligence (DFKI). He studied computer science and physics at the RWTH Aachen and received a PhD (Dr. rer. nat.) in computer science from Sarland University in 1995. His global research in computational linguistics has focused on linguistic formalisms, their efficient implementation and their mathematical foundations. He has worked on the morpho-syntax interface in HPSG, was the prime contributor of the typed description language TDL and the typed shallow processor SProUT, and led the group on deep linguistic processing in the Verbmobil project. His latest work in computational linguistics concerns the compilation of constraint-based grammar formalisms into weaker frameworks and the integration of deep and shallow processing methods. During the last few years, he has worked at the intersection of linguistic processing, general (world) knowledge representation and inference systems in the AIRFORCE, COLLATE and MUSING projects. He has implemented a time ontology that addresses temporal granularity and temporal underspecification in natural language and devised a general scheme for incorporating time into arbitrary description logic ontologies. His expertise in Semantic Web technologies includes theoretical frameworks, such as OWL, SWRL and SPARQL, as well as practical systems, like RACER, Pellet, OWLIM and Jena.

Markus Lankinen

Markus Lankinen, MSc (Economics), Vaasa, 1996, is Managing Director for the St-Petersburg-based consultancy company ManNet Partners and Regional Manager of the Finnish company FCG Planeko Ltd. As a management consultant, he is active in international business planning, specializing in start-up and development strategies. He is a member of the Finnish Association of Business School Graduates.

Christian Leibold

Christian Leibold holds a Diploma (master equivalent) of Computer Science from Munich Ludwig-Maximilians-University. The topic of his diploma thesis was the evaluation of content repository solutions and implementation at BMW Group Design Project Management. During his studies he concentrated on system architectures, specializing in Java EE and user interaction, including the teaching of tutorials accompanying the lectures in spoken dialogue systems and enterprise integration. He joined STI Innsbruck in April 2007. Since then he has won a research grant for the market preparation of a research prototype on the classification of information. In the EC IP MUSING he has a major role as project researcher. Before joining STI he was involved in industrial projects with BMW Group, Infineon Technologies, Take2 Interactive Software and Execcon Consulting, covering various parts of the engineering loop of software and hardware

products (e.g. requirements analysis, implementation, quality assurance, project management and process development).

Paolo Lombardi

Paolo Lombardi is Head of Internet Banking at Monte dei Paschi di Siena, the third largest Italian banking group. He joined the Montepaschi group in 2001, managing initiatives on financial-services-related innovation. From 2004, within the Organization Department, he undertook programme management responsibilities in the framework of the group's Basel II initiative. Since April 2007 he has also been responsible for the group's Contact Centre. He holds a *summa cum laude* Masters degree in nuclear engineering from the University of Pisa (1989), and has been appointed as Post-Graduate Researcher at the University of California Santa Barbara. He has been involved in the design and execution of European Commission research and technological development projects since 1992, and collaborates with the EC's DG Information Society as an evaluator and a reviewer.

Michael Munsch

Michael Munsch studied finance at the Comprehensive University of Essen. He received his PhD in the field of international finance risk management. Since August 2000 he has been Executive Manager of Creditreform Rating AG. He was previously Head of the Department of Risk Management at the Verband der Vereine Creditreform eV in Neuss, where he coordinated the development of new systems for credit rating and risk management. After finishing his studies in business administration he became a consultant for business clients at an international bank and was employed in the Department of Finance of a large international company. From 1989 to 1995 he worked in the Department of Finance and Financial Accounting at the University of Essen.

Richard Pike

Richard Pike, CCH SWORD Product Director, has more than 15 years' experience in risk management and treasury IT. He has analysed, designed and managed the development of core risk management systems for large international financial institutions. He was recently chosen as one of the 50 most influential people in operational risk, by *Operational Risk & Compliance* magazine. He is a regular speaker and writer on risk management issues.

Antonio Pippi

Antonio Pippi is a telecommunications engineer. In 2001–2005 he was with the Department of Information Engineering at the University of Siena, where he obtained a PhD with a dissertation on mathematical models for the study of