

# KNOWLEDGE AUTOMATION

How to Implement Decision Management in Business Processes

ALAN N. FISH

FOREWORD BY JAMES TAYLOR



### **Knowledge Automation**

How to Implement Decision Management in Business Processes

ALAN N. FISH





John Wiley & Sons, Inc.

Copyright © 2012 by Alan Fish. All rights reserved.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey. Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 646-8600, or on the Web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at http://www.wiley.com/go/permissions.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. No warranty may be created or extended by sales representatives or written sales materials. The advice and strategies contained herein may not be suitable for your situation. You should consult with a professional where appropriate. Neither the publisher nor author shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books. For more information about Wiley products, visit our web site at www.wiley.com.

FICO is a trademark of Fair Isaac Corporation in the United States and other countries.

#### Library of Congress Cataloging-in-Publication Data:

Fish, Alan N.

Knowledge automation : how to implement decision management in business processes / Alan N. Fish.

p. cm. — (Wiley corporate F&A series)

Includes index.

ISBN 978-1-118-09476-1 (cloth); ISBN 978-1-118-22351-2 (ebk);

ISBN 978-1-118-23679-6 (ebk); ISBN 978-1-118-26184-2 (ebk)

1. Information technology.—2. Decision making.—I. Title.

HD30.2.F574 2012

658.4'038011—dc23

2011042679

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

### **Foreword**

FIRST CAME ACROSS Alan's work in 2009 when he introduced me to the concepts at the heart of Decision Requirements Analysis. I was impressed by the approach he outlined and its potential for improving the analysis and design of decision management systems. I have been working in decision management for most of the last decade, spending much of that helping companies use business rules and predictive analytic technology to automate and improve business decisions. Alan's approach to gathering, modeling, and managing decision requirements immediately struck me as the right way to approach this problem. I have been using it with my clients ever since.

Decision management is a well established approach that focuses on automating and improving operational business decisions, including the many micro-decisions that impact a single customer or a single claim. In the years I have been working on decision management systems, the approach has become increasingly well known and adopted. Building decision management systems requires a solid platform for managing decision-making logic—a business rules management system—and the ability to integrate predictive analytic models with this logic. It also requires the ability to effectively identify, model, refine, and manage the requirements for such a system in a decision-centric way.

Take one company I was working with recently. They were struggling, trying to adopt a business rules management system into a process-centric culture. Using Alan's approach, we focused their energy on their decisions, on the dependencies between their decisions, and on the information and knowledge required by their decisions. Successful adoption of the business rules management system and a new appreciation for the role of decisions alongside processes were the result.

Many companies make the mistake of assuming that their existing requirements gathering and management processes will work well as they adopt business rules and a business rules management system. In fact, existing approaches lack the focus on decisions needed for success and offer little or no

help in actually analyzing decisions. Similarly, those using predictive analytic models often lack a formal or repeatable approach to defining the decisions their models are meant to assist. This is where Decision Requirements Analysis shows its worth.

This book is very timely. Decision management systems are increasingly topical and the focus of both the business rules and predictive analytics industries is shifting to a conversation about decisions and decision automation. In the past nine years I have helped dozens of companies that are implementing business rules and predictive analytics. I've spoken with literally hundreds more. Decision management systems are a powerful way for companies and organizations to improve their business and create a more agile, more analytic and more adaptive business. To succeed companies need more than great technology, they also need to be able to analyze and design these systems. Alan's work on Decision Requirements Analysis is a critical component of this.

With this book, Alan has written a practical and straightforward guide to using the Decision Requirements Analysis approach. A quick but thorough introduction to knowledge, decision-making, and the role of automated decision services in business processes sets the scene. The core chapter introduces the modeling approach and the workshop that goes with it. Alan shows how the approach helps at every stage of a project, from discovery to implementation, and outlines a great initial set of common decision patterns to help you get started. If you plan to analyze and model decisions—and you should—this book will show you how to do so.

I feel strongly that this is the best approach for modeling decisions out there, which is why I use it with my clients and recommend it in my most recent book.

—James Taylor

James is CEO and Principal Consultant, Decision Management Solutions, and is based in Palo Alto, CA. He is the author of Decision Management Systems: A Practical Guide to Using Business Rules and Predictive Analytics (IBM Press, 2011). He is an active consultant, working with companies all over the world, and can be reached at james@decisionmanagementsolutions.com.

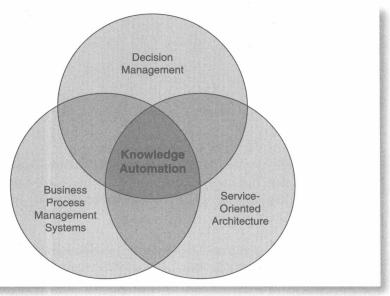
### **Preface**

RGANIZATIONS ARE CONSTANTLY searching for ways to make their business processes more effective, less costly, and more agile. Recently, the strategies adopted increasingly involve three technologies:

- 1. Decision management: the use of predictive analytics and business rules to optimize and automate decision-making, allowing day-to-day operational strategy to be controlled directly by process managers, rather than IT.
- 2. Business process management systems (BPMS): systems that allow business processes to be modeled as sequences of activities, and then ensure compliance in those processes by executing the models.
- 3. Service-oriented architecture (SOA): the provision of software functionality as loosely coupled, reusable services, which encapsulate their internal logic to hide it from the consumers.

This book addresses the intersection of these fields. Using a combination of all three technologies, it is possible to encapsulate business knowledge in services, consumed by the BPMS, which automatically make optimal operational decisions for the organization. This is knowledge automation (see the following graphic).

Process automation projects are always large and complex in any case, and usually represent a substantial strategic investment by the organization. After deploying any new IT infrastructure required, the organization faces the challenges of process redesign and organizational change. If operational decision-making is to be included in the scope of what is to be automated, we have an additional task: creating the services that are to make the decisions. Such "decision services" can be very complex, including sophisticated analytic models, hundreds or thousands of business rules, and elaborate calculations.



Knowledge Automation

Knowledge automation projects therefore need careful management if they are to be successful. Potential problems include:

- Development projects that deliver late or run significantly over budget
- Facilities that are not accepted by the users because they do not display a sufficient degree of competence
- Systems that are hard to maintain or get more confused year by year

These problems are often due simply to the lack of any clear definition of the requirements for decision making early in the project. In my experience, this is due never to a lack of effort, but rather to using the wrong approach. When these projects fail they have usually gone wrong at the very beginning.

There are a number of excellent books providing theoretical and technical background to decision management, business rules, and predictive analytics, and explaining why they are a good thing (a Suggested Reading list is provided in the back of the book). Although I include some basic background on the technologies and the benefits, this is not the main subject of this book. I assume you are already committed to or involved in a knowledge automation project, or at least have some understanding of the benefits of decision management and

are considering the feasibility or planning of such a project. The main purpose of this book is to give you a practical step-by-step guide to how to go about it.

Over the years I have been involved in a great many projects to build decision services for various business domains. Perhaps it is just my natural laziness, but I have developed a way of working that makes projects simpler and reduces the risks of failure and overrun. Simplicity is an underrated virtue: The methods we choose should be those that make tasks easier for people. People find a task easier when it has a clearly defined scope—that is, they know exactly what they have to do—and a clearly defined structure—that is, they can see how it can be broken into smaller tasks.

The method I describe in this book is Decision Requirements Analysis (DRA), which is used to define the scope and structure of the decision-making that is to be automated. DRA is quick and simple to use, but allows a domain of business decision-making to be defined with great clarity, because:

- It ties the decision-making to specific points in the business process: the decision points.
- It is based on a simple diagram of the structure of the decision-making: the Decision Requirements Diagram (DRD).
- It allows the scope of decision-making to be defined by listing the decision points, and drawing boundaries around parts of the DRD.

The primary statement of requirements in DRA is therefore a picture, a common language easily understandable by all involved: the project sponsors, project managers, process experts, business analysts, technical architects, and development team. The DRD allows the requirements for the automation of decision-making to be discussed and agreed by everyone at the start of a project. The structure it reveals then gives shape to the design, the implementation, and the project itself. I hope you will find it useful.



#### HOW TO READ THIS BOOK

Read it all. Go on, it's not long and not too technical—just read it all. Authors are expected nowadays to provide alternative paths through the text, tailored for different readers, but the approach I suggest depends on the idea that the structure of decision-making revealed in the DRD provides a framework for all the levels of a project, right from top to bottom, so to appreciate the value of the approach you need to see its full gamut. The structure of the book follows these levels:

Chapter 1 introduces the concept of the value of knowledge, first at the macroeconomic level with its contribution to economic growth, then in terms of its specific value to individual organizations. It provides a brief introduction to the principles of decision management, and describes decision yield, a way of measuring the benefits of automating decisions.

Chapter 2 looks at how business processes may be redesigned to automate operational decision-making, through the use of BPMS that call decision services at key points. It suggests ways of rationalizing the business process to derive maximum benefit from decision management, and provides a simple template originations flow as an example.

Chapter 3 provides some background on the most important technologies used for encapsulating business knowledge in decision services: business rules, algorithms, and predictive analytics, and explains when each should be used. It then describes how decision management involves a constant cycle of data through the BPMS, predictive analytics, and the decision services.

Chapter 4 presents the principles of Decision Requirements Analysis and introduces the Decision Requirements Diagram, which represents a domain of decision-making as a network of decisions, knowledge areas, and data areas. It describes in detail how to run a structured workshop resulting in the DRD and all the supporting information required to create an automation scoping document.

Chapter 5 shows how to use the structure of the DRD to scope, estimate, plan, and manage a project to implement a set of decision services, and how the same structure can be used as a basis for the design, development, configuration, and testing of those services, so that all activities are aligned efficiently to support a "knowledge production line," resulting in an implementation with full traceability of all components back to the original requirements.

Chapter 6 provides a set of commonly occurring decision patterns, useful as templates, with simplified partial DRDs and object models. Finally, it suggests a number of alternative ways of handling decision-making where decision services and business users must collaborate on reaching a decision, and suggest situations in which these might be used.



#### **INTENDED READERSHIP**

A knowledge automation project involves the participation of many people in a great cascade of activity. This book is intended to help them all.

- Executives in organizations considering or implementing automated decision-making, especially the CEO, CIO, and CFO, who need to understand the basic principles of the technologies and the issues involved in implementing them
- IT architects involved in the procurement of BPMS, BRMS, and analytic products, who need to know the functionality required for knowledge automation and how the products should interoperate
- Analysts working in the design or redesign of business processes to use automated decision-making, who need an approach that will result in a clean and efficient process design
- Change managers and personnel managers, who need to understand how the roles of staff change when knowledge automation is introduced and how to be involved in defining that change
- Managers responsible for aspects of organizational performance to be addressed by the automation (e.g., product managers, risk managers, and process managers), who need to ensure that the new systems will support them in more agile and effective management
- Project managers, who need to know how a knowledge automation project should be structured and managed so as to minimize costs, timescales, and risks
- IT architects involved in the design of systems to automate decisionmaking, who need a set of principles for specifying a set of decision services and designing their internal structure and interface signatures
- Analysts responsible for defining the business requirements for automated decision-making, who need a top-down approach for discovering and codifying business knowledge

I hope the book will also be of some interest to those whose concerns are more theoretical than practical.

- Standards bodies and consortia, which are interested in defining executable models of decision-making encompassing multiple forms of business knowledge
- Academic institutions running courses in IT and AI, which wish to provide students with a robust methodology for knowledge automation
- General readers, who are interested in what happens behind the scenes when they apply for a financial product and how institutions make automatic decisions that affect them directly



This is a "how to do it" guide, not a theoretical treatise, intended principally to be used by those actually involved in knowledge automation projects. It must therefore be directed at the prevailing paradigm for process automation: BPMS and SOA. In particular (this will be discussed in detail in Chapter 2), it assumes that stateless decision services will be called from process flows managed by a BPMS. So, for example, it neglects the emerging field of complex event processing (CEP), in which decision-making is triggered by the detection of an event, rather than by a task in a process flow. However, the central concepts in this book—concerning how to represent and implement the structure of automated decision-making—are applicable however that decision-making is triggered.

The book aims to be product-neutral throughout and therefore restricts itself mainly to functional issues, rather than technical ones. You will find here no comparison of the facilities of different BRMS products, no explanation of how inference actually works in a BRMS, and no discussion of the systems architectures required to support BPMS and BRMS. As a result, I hope the methods and techniques I describe in this book will be applicable whichever products and architectures your organization elects to use.

As suggested by the title, the book focuses on the automation of operational decision-making in business processes, rather than any other uses of decision management technologies (e.g., product design and marketing, business intelligence, and data mining). It also uses examples taken from mainstream application areas of process automation (to be frank, those that are most familiar to me). These include originations, account management, and claims management in banking, retail credit, general insurance, life insurance, and health insurance. Of course, knowledge automation can be applied in many other scenarios.

### **Acknowledgments**

owe a profound debt of thanks to many people:

- My wife Cath for her unswerving support while I worked on this book, and my younger children Oscar and Penny, who have had to tolerate my frequent absences.
- My colleagues at FICO, for indulging me in the implementation of these ideas, helping to promote them in the industry, and being such good company when working on-site.
- All the clients on many projects, from whom I have learned so much, for giving me their trust and hospitality and making this work so thoroughly enjoyable.
- My friend Roger Farmer, for useful advice, a shot of confidence, and too many shots of tequila.
- James Taylor, for his support, his fount of knowledge, and for kindly agreeing to provide the Foreword.
- Paul Konnersman, for engaging in lengthy correspondence on approaches to decision automation, and allowing me to cite his work, even though he doesn't entirely agree with me.
- My late father, for teaching me how to think about systems. I miss arguing with him.

I am sketching this introduction in the brick-arched cellars of Jazzland, the oldest jazz club in Vienna. I'm eating cevapcici; drinking Austrian beer; and watching Vincent Herring on sax, Danny Grisset on piano, Markus Schieferdecker on bass, and Joris Dudli on drums: all astonishing musicians and a seriously cool crew. I must confess, I am extraordinarily privileged to be able to work in an interesting field, doing a job I love, with intelligent clients and colleagues, in beautiful places. Long may it continue.

### Contents

vii

Acknowledgments XVII	
Chapter 1: The Value of Knowledge	
The Economics of Knowledge The Knowledgeable Business Notes	14
Chapter 2: Decisions in the Business Process	15
Business Process Modeling Decision Points and Decision Services Redesigning Process Decisions Originations Process Template Approaches to Process Design Notes	16 17 20 33 36 39
Chapter 3: Encapsulating Knowledge in Decision Services	41
Business Rules Algorithms Predictive Analytics Putting It All Together Notes	43 52 56 65 68
Chapter 4: Decision Requirements Analysis	71
Principles  Decision Requirements Diagram	72 73

Foreword ix

Preface xi

#### viii Contents

DRAW Notes	79 91
Chapter 5: Using DRA in Knowledge Automation Projects	93
DRA in Project Management	94
DRA in Knowledge Discovery	107
DRA in Design	115
DRA in Implementation	124
The Knowledge Production Line	131
Notes	133
Chapter 6: Common Decision Patterns	135
Implementation Types	135
Collaborative Decision-Making	147

Conclusion 155

Glossary 159

Suggested Reading 167

About the Author 171

Index 173

CHAPTER ONE

## The Value of Knowledge

HAT COMES INTO YOUR HEAD when you hear the word "knowledge"? Dusty books in the libraries of ancient universities, perhaps, or men in white coats with big foreheads? My aim in this first chapter is to change the way you think about knowledge. It is not an abstract concept divorced from the world of business; it is a tangible corporate asset. You can manufacture it, own it, buy and sell it, build it into machines that make profits for you: It is real *stuff* that has *value*. This is a crucial concept in understanding the potential benefits of decision management.

I'm going to start right at the top with a brief discussion of modern macroeconomic theories of knowledge before looking more practically at how organizations can measure and exploit the value of their own business knowledge.



#### THE ECONOMICS OF KNOWLEDGE

The past 50 years have seen a revolution in our understanding of the drivers of economic growth. Classical economic theory addressed labor, land, and capital. In much the same way, modern economic theory now addresses human knowledge, and its contribution to the growth of economies.

### **Neoclassical Growth Theory**

Theories of growth proposed in the 1950s, such as those of Robert Solow<sup>1</sup> and Trevor Swan<sup>2</sup> were based squarely on physical capital: objects such as machinery and stock. These were called "neoclassical growth models." According to neoclassical theory, physical capital is assumed to be subject to the law of diminishing returns: the principle that if you keep increasing any single factor of production, the output you achieve from each unit of input will eventually decrease. This assumption is mathematically convenient, because it results in a model economy that always converges to a unique steady state. In other words, under any constant conditions, growth will slow down and eventually stop.

Figure 1.1 shows the attitude prevailing at the time. Bill Phillips built this hydraulic computer—MONIAC—in 1949, using water flowing around pipes and tanks to simulate the U.K. economy. This was a sophisticated and accurate model, but the message is very clear: the economy finds its own level.

Neoclassical models could not explain long-term growth, and were not intended to. In this approach, long-term growth is accounted for by external influences. Robert Solow introduced technical knowledge as such an "exogenous variable." <sup>3</sup> He assumed that technology was improving steadily without any influence from the economy being modeled; it simply happened, providing the external stimulus that kept the economy growing. Such a model allows you to measure the *effects* of technological progress, or to decide, for example, the optimal savings rate given a certain rate of progress. Because of this, Solow's model became widely used in economic analysis and earned him the 1987 Nobel Prize for Economics. What his model cannot do, however, is help you to determine economic policy to *achieve* technological progress.

### New Growth Theory

The thrust of more recent theories of growth has therefore been to "endogenize" knowledge: to bring it within the terms of the model as an internal variable and use it to explain the observed growth of economies. Robert Lucas explains why such models are so important:

Is there some action a government of India could take that would lead the Indian economy to grow like Indonesia's or Egypt's? If so, what, exactly? If not, what is it about the "nature of India" that makes it so? The consequences for human welfare involved in questions like these are simply staggering: Once one starts to think about them, it is hard to think about anything else.

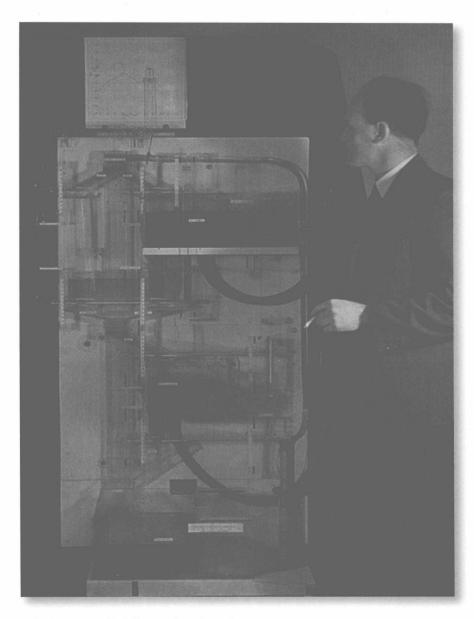


FIGURE 1.1 Bill Phillips with MONIAC