



The Decision Model: A Business Logic Framework Linking Business and Technology

Barbara von Halle and Larry Goldberg

CRC Press

\$52.22 521 pages

Reviewed by Paul Harmon

One of the most exciting developments in Business Process Management today is the gradual merger of business process analysis and business rules analysis. Only a few years ago the two groups hardly talked with each other. Today it is widely recognized that process people benefit from describing decisions by means of business rules and the rules people benefit by having a well-thought out process structure to use in defining where business rules are used.

In fact, the growing synthesis has many dimensions. For example, there is a move within the rules community to broaden their concerns from rules to decisions, allowing the same analysis effort to define either decision tables or business rules, depending on what best fits the business need. In a similar way, checklists can be seen as a kind of decision device.

My own background includes a decade in Artificial Intelligence (AI) in the Eighties, and much of the recent work reminds me of the interest in expert systems in that era. The first expert system tools focused on defining rules – and often produced hundreds or thousands of rules to be used in reaching a decision. At some point, the problem of maintaining the large rules sets became so onerous, that theorists began to combine frame-based systems with rules. Frames were simply an early, AI-derived name for what soon became objects. By the mid-Eighties, all the leading expert system tools and methodologies supported both rules and frames. In essence, one created an object hierarchy to define a subject matter, and then used rules to reason about occurrences. Thus, if one wanted to identify an animal, one started with a hierarchy in which the top object was Animal. Subclasses might include Birds, Mammals, Fish, and Invertebrates. The Animal class described the characteristics that were common to all animals – breathed, moved, was born and died, etc. The Bird class “refined” the Animal class with additional features characteristic of Birds, including: has Feathers, lays Eggs, and has Wings.

Rules were associated with each class to reason about specific features. Thus, one began by determining if the object in question was an animal, and then applied rules to determine if the object in question was a bird or a fish, or something else.

Much of what can be done with object hierarchy can be done with a relational database, although it is harder to do.

The other key difference between most expert systems and much recent business rules research was the nature of the knowledge that was captured in the expert system. As a broad generalization, business rule systems have focused on capturing business policies and their implications. Expert systems were more likely to focus on capturing the knowledge on an individual who solved problems as a result of accumulating many years of experiential knowledge. Thus, an expert system application was built to mimic the work of a petroleum engineer who could reliably identify sites for oil drilling, or a physician who was a world expert on the diagnosis of meningitis infections.

Obviously it is easier to capture policies and their implications than the tacit knowledge of human experts. One begins with an explicit and logical statement of a company's intent. The other begins by sitting with an expert and asking how they would approach a specific problem, and then repeating the process over and over again until one begins to have an understanding of the expert's reasoning process.

Process practitioners face both kinds of problems. In many cases companies want help assuring that company policies are fully and correctly implicated. A rule based system based on the companies policies may serve very well. In other cases, a process team wants to capture the knowledge of a specific individual and make that individual's experiential knowledge more widely available.

In the late Nineties and early Zeros, it seemed as if Rules theorists were determined to focus on the relatively simple problems involved in that capture of policy information. These problems are complex enough, requiring as they do, consistency and discipline in the use of terms throughout the organization. But they usually result in relatively modest rule sets that can be manipulated by relational databases. In the past few years, however, business rules theorists and decision theorists have begun to reconsider the problems and the techniques initially explored by the expert systems designers in the Eighties.

Barbara von Halle is a well-establish business rules guru. She and her business partner, Larry Goldberg have written a book that will challenge most of today's business rules groups. In essence, they are concerned with how to manage larger rule bases, and argue that the rule sets need to be divided into sets and that the sets need to be structured by tables that identify the key terms used in each set. They propose, in other words, to resurrect the "frames and rules" approach used by the expert systems developers of the Eighties. They are still primarily focused on relatively simple policy-derived rule sets, so relational databases provide them sufficient structure for their purposes, although if this work were to be carried much further, users would find one of the Object-Oriented Databases (which were originally created to support expert systems work) much easier to use.

Using the idea of structuring the knowledge to be used into a hierarchy, von Halle and Goldberg show how it is easier to create and much easier to maintain large rule sets. At the same time, they lay a foundation for linking rules and decision table approaches, laying the foundation for what I assume will be the next generation of business rule/decision systems.

This is not a book for a business manager. This is a book for someone charged with understanding how to create business rule or decision support systems. That said, *The Decision Model* is an important book for a technically sophisticated audience. It should be read by those engaged in business rules work to learn about better ways of organizing rule sets. And it should be read by process practitioners to learn about the growing merger of decision and process techniques. To help assure that the book will prove relevant for process people, von Halle and Goldberg have asked other authors to write key chapters, and thus John Zachman, James Taylor, Mike Rosen, Bruce Silver, and Daniel Worden have contributed chapters, and they have been supplemented by authors from institutions that have conducted experiments with the von Halle and Goldberg methodology.

The book is important because it introduces better ways to manage large rules bases and more flexible ways to think about decisions in the context of processes. As an interest in dealing with complex, dynamic processes (case management) grows, however, there will also be a growing interest in learning how to capture tacit knowledge that knowledge workers have acquired, and I confidently predict that there will be a growing interest in small “decision support” systems that follow along the lines of earlier expert systems. These systems will require techniques that von Halle and Goldberg discuss in this book, and I believe that many rules and process practitioners will find this book an important step toward coming up to speed on the new techniques they will want to use in the course of the coming decade.

Paul Harmon is the executive editor of www.bptrends.com

BPTrends LinkedIn Discussion Group

We recently created a BPTrends Discussion Group on LinkedIn to allow our members, readers and friends to freely exchange ideas on a wide variety of BPM related topics. We encourage you to initiate a new discussion on this publication or on other BPM related topics of interest to you, or to contribute to existing discussions. Go to LinkedIn and join the **BPTrends Discussion Group**.