# Introduction

# 1 What is this book about and who should read it?

This book is about facilitating change with component technology, but is different from most approaches to the topic. The components in the book are not traditional I/T components. Rather they are shared components of knowledge from which patterns of business knowledge are assembled. A fundamental premise of this paradigm is that meanings are patterns of information, abstract structures that may be derived from other components, which are also meanings. If we can identify and describe these components and their structures with precision, we can automate the process. Business processes and information systems configured from these components will be extremely flexible, configurable, and coordinated.

This book lays the foundation of a new computing paradigm – a paradigm in which computers manipulate meanings, not program code or blind symbols. Computers of the future, built on the principles described in the book, will operate on the plane of meanings – a little like we do.

Business meanings, patterns, and rules jointly constitute the substance of a business process. Without the business layer, technological standards have little meaning. The return on investment from reusing business knowledge can complement, and be orders of magnitude larger than by adherence to technical standards alone. This book establishes a framework for the transfer and reuse of business knowledge in different contexts. This is why we urge architects interested in service oriented architecture and business process management to read this book. This book is for information and process architects. It is also a book for architects of *languages* for specifying business processes (languages like BPMN and information sharing standards such as SBVR and BPDM from object Management Group (OMG).).<sup>1</sup>

This book is about automating the configuration of business processes from components of business knowledge. We urge software architects and technologists to read this book

1

<sup>&</sup>lt;sup>1</sup> The Business Process Management Initiative (BPMI) consortium is a consortium of diverse firms. Its purpose is to standardize business process definitions "that span multiple applications, corporate departments and business partners, behind the firewall, and over the Internet." BPMI has published the BPML language in support of business process automation. See http://www.bpmi.org. XML is from W3C, another consortium for data standards.

2

Cambridge University Press 978-0-521-85163-3 — Creating Agile Business Systems with Reusable Knowledge A. Mitra , A. Gupta Excerpt More Information

#### Creating Agile Business Systems

because it is about the technical principles that drive information architecture and information sharing in the form of meanings and concepts. The principles and patterns in this book complement the work that has been already been done in developing technology and interfacing standards for information systems. The purpose of this book is not to propose yet another technical standard. It is to describe the business intelligence, in component form, that these standards must support and be joined to. It is the next step.

# 2 What will the information be used for?

The patterns in this book can address the following business issues:

- 1 Agility and adaptability of businesses processes and systems: facilitate designing of agile business processes and flexible systems based on the reusable patterns of information in this book. They will help automate the alignment between business processes and information systems, speed development of systems to support new products and distribution channels, and accelerate process and systems integration when businesses integrate or reinvent themselves in their product markets.
- 2 *Integration and coordination of information*: coordinate integration of information and processes across supply chains, enterprises, and databases.
- 3 *Reduction of time-to-market of new concepts*: accelerate formulation of functional requirements and process models based on the prefabricated reusable patterns in this book.
- 4 *Creation of automated tools for aligning of information systems with business:* facilitate the development of an integrated Computer Aided Process Engineering (CAPE) and Computer Aided Systems Engineering (CASE) tool; provide a framework for testing the completeness and validity of a language or methodology for business rule/process definition and modeling, and for evaluating CAPE and upper CASE tools.
- 5 *Compression of the time to develop prototypes*: the patterns delineated in this book can serve as the basis for early prototypes when iterative prototyping methodologies are used for developing or integrating information systems and business processes.

Ultimately this book is about change. It describes a technology for automating and facilitating change – a technology that will facilitate the innovation and adaptation so necessary for corporations to remain competitive in a fast-changing, diverse, and tumultuous business universe that will not forgive the tardy.

# 3 Technology's broken promise

Why is change so difficult? That is a question with an easy answer. We have all experienced how a seemingly simple change to a business process or information system has many impacts – often unanticipated, at multiple places, in multiple ways – on other different business processes and at many different layers of the information systems legacy that support these processes. Each impact has several other impacts in turn, which ricochet through our processes and systems until we are caught in an explosive cascade of change. Business sponsors requesting the change are then faced with a painful choice – either make



Figure 1 Information is a key resource, but investment in information systems is fraught with risk

the changes at a cost, both in time and money, and take a risk that might be excessive, or abandon the competitive benefits of innovation because the risk is too high and the change is not timely. Studies have indicated that information system projects are fraught with risk, and businesses do not realize the value they should from their investments in information systems (figure 1).

A recent example was the Y2K problem that seemed trivial to the layman, but the stateof-the-art of computer technology was such that it may have cost industry as much as \$600 billion<sup>2</sup> and a significant part of the world's resource pool of professionals to merely express the year in four, instead of two, digits.<sup>3</sup> Many strategic benefits have been difficult to implement for similar reasons. Denial of strategic benefits to consumer and provider alike are so frequent that examples litter the industrial landscape in almost every direction. Examples of missing capabilities include:

- Straight through processing and "T+0" settlement in the financial securities industry (the ability to settle a trade immediately with almost no manual processing).
- Real time billing for telephone subscribers and personal telephone numbers in the telecommunications industry (where a unique contact phone number automatically follows an individual regardless of location or geography).
- Timely and reliable order fulfillment and innovative customer service for manufacturers and retailers.
- Risk assessment when providing insurance coverage to complex global clients in the insurance industry.

This is only a small slice of such wish lists – strategic innovations and improvements in almost every industry are often deemed too risky or impossible because supporting processes and information systems are deemed too complex and risky, if not impossible, to change.

Despite inventing new technology at a prodigious rate to make change easy (including technologies such as CASE tools, code generators, structured programming, relational databases, expert systems, object technology, and reusable components), systems still cannot

<sup>&</sup>lt;sup>2</sup> Sources: Gartner Group & Congressional Research Service estimates quoted by Steve C. Yuen, Ph.D., University of Southern Mississippi and Jo Ann Mitchell, Jones Junior College on our website.

<sup>&</sup>lt;sup>3</sup> For example, 1/1/2001 instead of 1/1/01. Computer calculations involving dates beyond 1999 had a very high risk of error if the year was not expressed in four digits.

#### Creating Agile Business Systems

change fast enough. Why are information systems a bottleneck? Can process re-engineering, business innovation, and time to market be accelerated? Why have these technologies not fulfilled their full potential?

The principal reason why the problem of change continues to persist is that we have not found a way of representing business rules and knowledge in a single place in such a way that we can change a rule once, and reflect the change wherever it impacts business processes and supporting automation. Instead, rules of business are repeated in different forms and formats in multiple, tangled ways in several systems at several places, which makes change complex, risky, and difficult. This has been a core problem.

It was not solved in the 1950s when we replaced the tangled code of machine language with assembler languages, or in the 1960s when we replaced the spaghetti code of assembly languages with that of third-generation languages like COBOL and FORTRAN; nor was it solved in the 1970s and 1980s with the coming of relational databases, expert systems, and CASE tools, or even in the 1990s when tangled object inheritance became so much of a problem that many advocated making multiple inheritance illegal in tools of the day. *More automation merely automated tangling of more business rules faster*.

For this reason, the authors asked a different question: what information do we need to model the stimulus response behavior of business processes and the organizations they support in the real world, and what is the *natural* real-world structure of information that can represent business knowledge in fully normalized,<sup>4</sup> and hence reusable form across the universe of diverse global business environments?

Why would the proposed approach work when so many others have failed? It works because it untangles business rules. It untangles business rules even if they were tangled in legacy models and systems. Thus it allows us to represent business knowledge *once* in a repository of knowledge, from where it can *naturally* manifest itself in different business contexts. Changes made in the right place will automatically impact business systems where they must. It is no surprise that many businessmen and professionals have intuitively felt that business knowledge acquired in one context might often be reused in another. We discovered in 1992 that this intuition is a fundamental truth that flows from the natural structure of business knowledge in the real world. However, we must explicitly recognize this structure and express it with mathematical precision to use it effectively. We will share this vision with you in the chapters that follow.

## 4 Component reuse – the genesis

The concept of using reusable components to compress application development time is almost as old as the software industry. Components have evolved from concepts such as copy libraries, common subroutines, and general purpose applications packages, in the early days of batch computing, to reusable GUI, network, and data services objects, based on standards such as CORBA, COM, and web services such as XML and WSOL which support distributed, interactive Web, and client–server computing.

<sup>4</sup> "Normalized" means represented uniquely in a single place once.

#### 5 Introduction

For historical reasons, software component engineering first focused on the back end of the process engineering value chain. Its first concern was program code and interfaces for communicating data in terms of streams of bits and bytes. The economic impact, however, is usually far larger at the front end of the re-engineering value chain – on reusing business knowledge to configure and innovate business processes, services, and products. It is not surprising that business had only very limited interest and no involvement in the kind of components that software engineers were interested in. Consequently, the business community's support for the software community's component technology was lukewarm at best. The focus has shifted in step with evolving technology. Now the time is ripe to look at the reuse of knowledge. This territory, long neglected by the software community is, and has always been, where the major benefits to business are found. Let us analyze these imperatives.



The supply chain for IT solutions represents the process of technology initiatives, application development, implementation and business use. For example, the supply chain for an ERP system may comprise of technology platform selection, systems specifications, systems development, packaging and documentation, implementation, and use in the end user business. The concept of the demand chain, which transfers demand from end users to technology suppliers, is less familiar. To give one example, the demand chain for an ERP may start with business users spotting new opportunities for using the system to support their business. The next link in the chain is the IT

6

Cambridge University Press 978-0-521-85163-3 — Creating Agile Business Systems with Reusable Knowledge A. Mitra , A. Gupta Excerpt More Information

#### Creating Agile Business Systems

department of the user organization looking for potential solutions already in place in the business. In the demand chain of the ERP system, it is 'missing' process solutions that drives the next stage – a process innovation stage – where new processes and solutions are outlined. The last step in the demand chain, is demand for resources and skills needed for using, operating, and developing the ERP system in the user organization over time . . . What is needed is capabilities to capture an increasing number of business opportunities already in the use . . . the supply chain for IT solutions needs to be managed so that both the current and future applications architectures are scalable, flexible and modular. (Jan Holmström and Tiina Tissari, *IT Value Capture: Creating an Effective Demand–Supply Chain for IT Solutions*)

As businesses have become increasingly reliant on automation, the line between technology and the corporation's key business operations has started to blur. Industry has begun to recognize that the greatest benefit to business will flow from reusing business intelligence embedded in software. Consequently the software industry has been striving to craft software components to reuse this embedded business intelligence across the supply chain. The intent is to speed up business processes, to make corporations more agile, and, above all, to position the business at a competitive advantage.

However, this kind of reuse has remained elusive in spite of over 15 years of industry effort. The reason why the promise has not been fulfilled is that industry was not ready to leverage the technology – processes had not matured, technology was still groping for the right answers, software developers were loth to frontload effort on software projects, and, most of all, business sponsorship was weak because software architecture was not as critical to successful business as it is today.

E- (as well as M-) commerce has forced cross-enterprise transparency into business processes and driven the need for standards. The market is now ripe for a product offering software components that will encapsulate and reuse business knowledge to build software architected to facilitate business innovation, speed, and agility: software that must be developed in compressed timeframes. Competitors are few and it is a prime opportunity for entrepreneurial corporations willing to take the plunge to build and sell software components that reuse business intelligence.

#### THE NEW OPPORTUNITY

Components will reduce the need for large scale integration – the bread and butter of the Big Six and others. But even as they mourn the loss, a new practice will emerge: helping users pick from the rapidly growing set of component based options. (The Forrester Report, "Package Application Strategies," June 1, 1996)

## 5 Scope of this book

In the following chapters, we will examine how project managers, requirements analysts, process engineers, and information modelers can leverage the frameworks and patterns proposed in this book to do their jobs faster, better, and with fewer resources.

Balancing risk with reward is at the heart of every business. Therefore it might be ironic, but hardly surprising, that business operations are largely deterministic. They are designed to minimize uncertainty. The scope of this book is therefore limited to deterministic patterns

#### 7 Introduction

and processes. This assumption will simplify our model of knowledge. However, we cannot do justice to business knowledge if we ignore uncertainty and risk altogether.<sup>5</sup> Therefore the model of knowledge in this book does provide some structures and components that partially, but adequately,<sup>6</sup> compensate for the purely deterministic nature of the model.

Unlike engineering systems, the vast majority of business processes deal with discrete, not continuous, change. For example, a business agreement might be negotiated in discrete steps that start with a first draft, progress through a series of reviews and revisions, and finally end with a binding contract. The scope of this book is limited to models of discrete, not continuous, change. This is adequate for almost all business processes. An engineering process, however, might contend with feeds and parameters that change or flow continuously. For example, when hydrocarbon-based resins are made in a chemical reactor, the density of the resin produced varies continuously with changing temperatures and pressures in the reactor. Instead of focusing on continuous technical processes, our focus, is on the discrete business process.

The book focuses on normalizing,<sup>7</sup> encapsulating, and reusing business knowledge across the value chain described in box 1. Business knowledge is technology independent. This knowledge may be embedded in processes that are supported by diverse technologies, both automated and manual. Often, in large organizations, the same business rules are expressed in different systems and procedures, on different technology platforms, in different countries or organizational units. The choice of the technology often depends on the organization's legacy, its local environment, and its infrastructure. Although business knowledge is independent of the technology that implements it, if an organization wants to reuse business knowledge in such a repository is an item of information that is expressed in some physical format and medium and is an electronic artifact. For this reason, we have named these components of knowledge *business knowledge artifacts*.

*Business knowledge artifacts* complement, but are different from traditional software components. The following chapters will show you how to link business knowledge artifacts to software components.

Because this book focuses on the rules of business, *business knowledge artifact* has often been abbreviated to *knowledge artifact* in the material that follows. *Knowledge artifacts* encapsulate bits of formal business intelligence – meanings – that can be stored as reusable components in a repository of business knowledge. Standardized knowledge artifacts will be central to the evolving knowledge economy, especially to the global supply chains emerging in the new economy.

## 6 Foundation of knowledge reuse: three pillars

Business knowledge is not about files, data flows, formats, screens, or computers. Rather it is about processes, practices, norms, products, policies, regulations, infrastructure, and people, constrained by the physical, regulatory, and ethical contexts in which they function.

<sup>&</sup>lt;sup>5</sup> Publications in the bibliography discuss stochastic processes (processes based on chance and uncertainty).

<sup>&</sup>lt;sup>6</sup> Adequate for the vast majority of business processes. <sup>7</sup> See the endnote on normalization.



Figure 2 Reusing business knowledge: the three pillars

To recast this knowledge in the form of normalized<sup>8</sup> and reusable capsules of information that can be assembled into configurations of knowledge and innovative ideas, we must first know how knowledge can be normalized. We must also know which parts are reusable and how to organize people and business practices to leverage these reusable knowledge components. The three pillars in figure 2 are the pillars on which business knowledge components must stand.

## 6.1 The first pillar: metamodel of business knowledge

Knowledge can only be reused if it is extracted and stored as a single piece of information in the knowledge repository. This information can then be used in as many different contexts as necessary, whenever and wherever it is needed. Additionally, in order to track its impact, we must know the relationship this piece of information has with other similar bits of knowledge in our repository of business knowledge.

<sup>8</sup> Normalization is a structured method of representing information in a non-redundant way. The endnote on normalization describes it in more detail.

#### 9 Introduction

For example, if the exchange rate between the US dollar and British pound changes, it could impact several valuations such as amounts on invoices, credit limits, checks, payment amounts, cash on hand, and fixed assets overseas. In other words, we must know the *structure* of information in the real world – that there are interrelated entities such as processes, resources, work products, and units of measure.

This information about information is collectively called a *metamodel*. The metamodel will provide the scheme for storing business knowledge in a non-overlapping, non-redundant way. The abstract objects in the metamodel, such as process, resource, unit-of-measure, and their interrelationships, are containers of non-redundant (normalized) business knowledge. Individual business knowledge artifacts would be classified and stored in these containers provided in the knowledge repository.

Specific knowledge artifacts can then be extracted from these containers and assembled into complex business processes and bodies of knowledge around which information systems can be built, the metamodel is the schema of the knowledge repository. It is the first pillar on which knowledge reuse stands. Without it, there can be no knowledge components. This book develops the metamodel of business knowledge. A companion book *Knowledge Reuse and Agile Processes, Catalysts for Innovation* extends the metamodel.

#### 6.2 The second pillar: business patterns

How many business rules does an enterprise need in order to do business? We know that only a small fraction of business knowledge is explicitly recorded and recognized by most operating businesses. Most business knowledge is implicit. Some are common sense rules that seem foolish to explicitly publish, such as "accept payment for goods sold," while others might be embedded in the experience or common understanding of the firm's employees, such as "breaking my budget will be a career limiting move." However, automation has no innate commonsense unless it is explicitly built in. Extracting and storing all rules of business, implicit and explicit, for even a small and simple business like a mom and pop corner store is not just a daunting task, it is an impossible task (figure 3): there are too many rules. There can be only one outcome if an analyst attempts to discover *every* rule of business for even the simplest enterprise: analysis paralysis.

Fortunately there is a solution. The knowledge repository is an electronic warehouse that holds the inventory of knowledge components and facts about how the business operates. Manufacturers and retailers who deal with large and diverse component and product inventories stored in brick and mortar warehouses are familiar with two fundamental laws of inventory management:

- 1 Only a few kinds of items account for the most frequent movement of inventory. Businesses need the vast majority of other items less frequently.<sup>9</sup>
- 2 Only a few items (not necessarily those with volatile inventories) are most critical to the business.<sup>10</sup>

<sup>&</sup>lt;sup>9</sup> The analysis of which items in inventory are needed frequently and which infrequently is called ABC analysis. Category A item inventories are the most volatile and category C the least volatile.

<sup>&</sup>lt;sup>10</sup> An analysis of the degree of criticality of items in inventory to business operations is called 123 analysis. Category 1 items are the most critical to business and category 3 items are the least important.



**Figure 3** Analysis paralysis: only a few critical rules, reused often, connect the business of the enterprise but they are lost in a tangled web of minutiae

Knowledge inventories too follow these laws. Every rule of business need not be extracted and stored before the business can benefit from the concept of knowledge artifacts. Also there are only a few critical items of business knowledge that are reused most often. These items can be discovered in common business patterns that not only orchestrate the internal operations of the enterprise and its many diverse functions, but also connect the enterprise to stakeholders across supply and demand chains. This is also the knowledge that is of utmost value to the business and impacts it maximally! (See box 1.)

One of the authors served as the director of systems architecture for NYNEX at the time when it was one of US's largest telecommunications firms. NYNEX was then wrestling with the implications of the impending deregulation of the US telecom industry. Earlier, when he had worked for AIG, a large insurance firm, he had identified several fundamental patterns of business that were common to all businesses, regardless of what they produced, or where they were located. He was delighted to find that these common patterns (from his AIG experience) could be applied to the core processes, products, and services of the telecommunications industry as well. Indeed, they even anticipated key changes driven by deregulation before users articulated their requirements. We believe that the opportunity to leverage knowledge patterns and artifacts exists for many other industries and applications, and show how this may be done in the chapters that follow in this book and in the modules available on our website.