# IDIOM

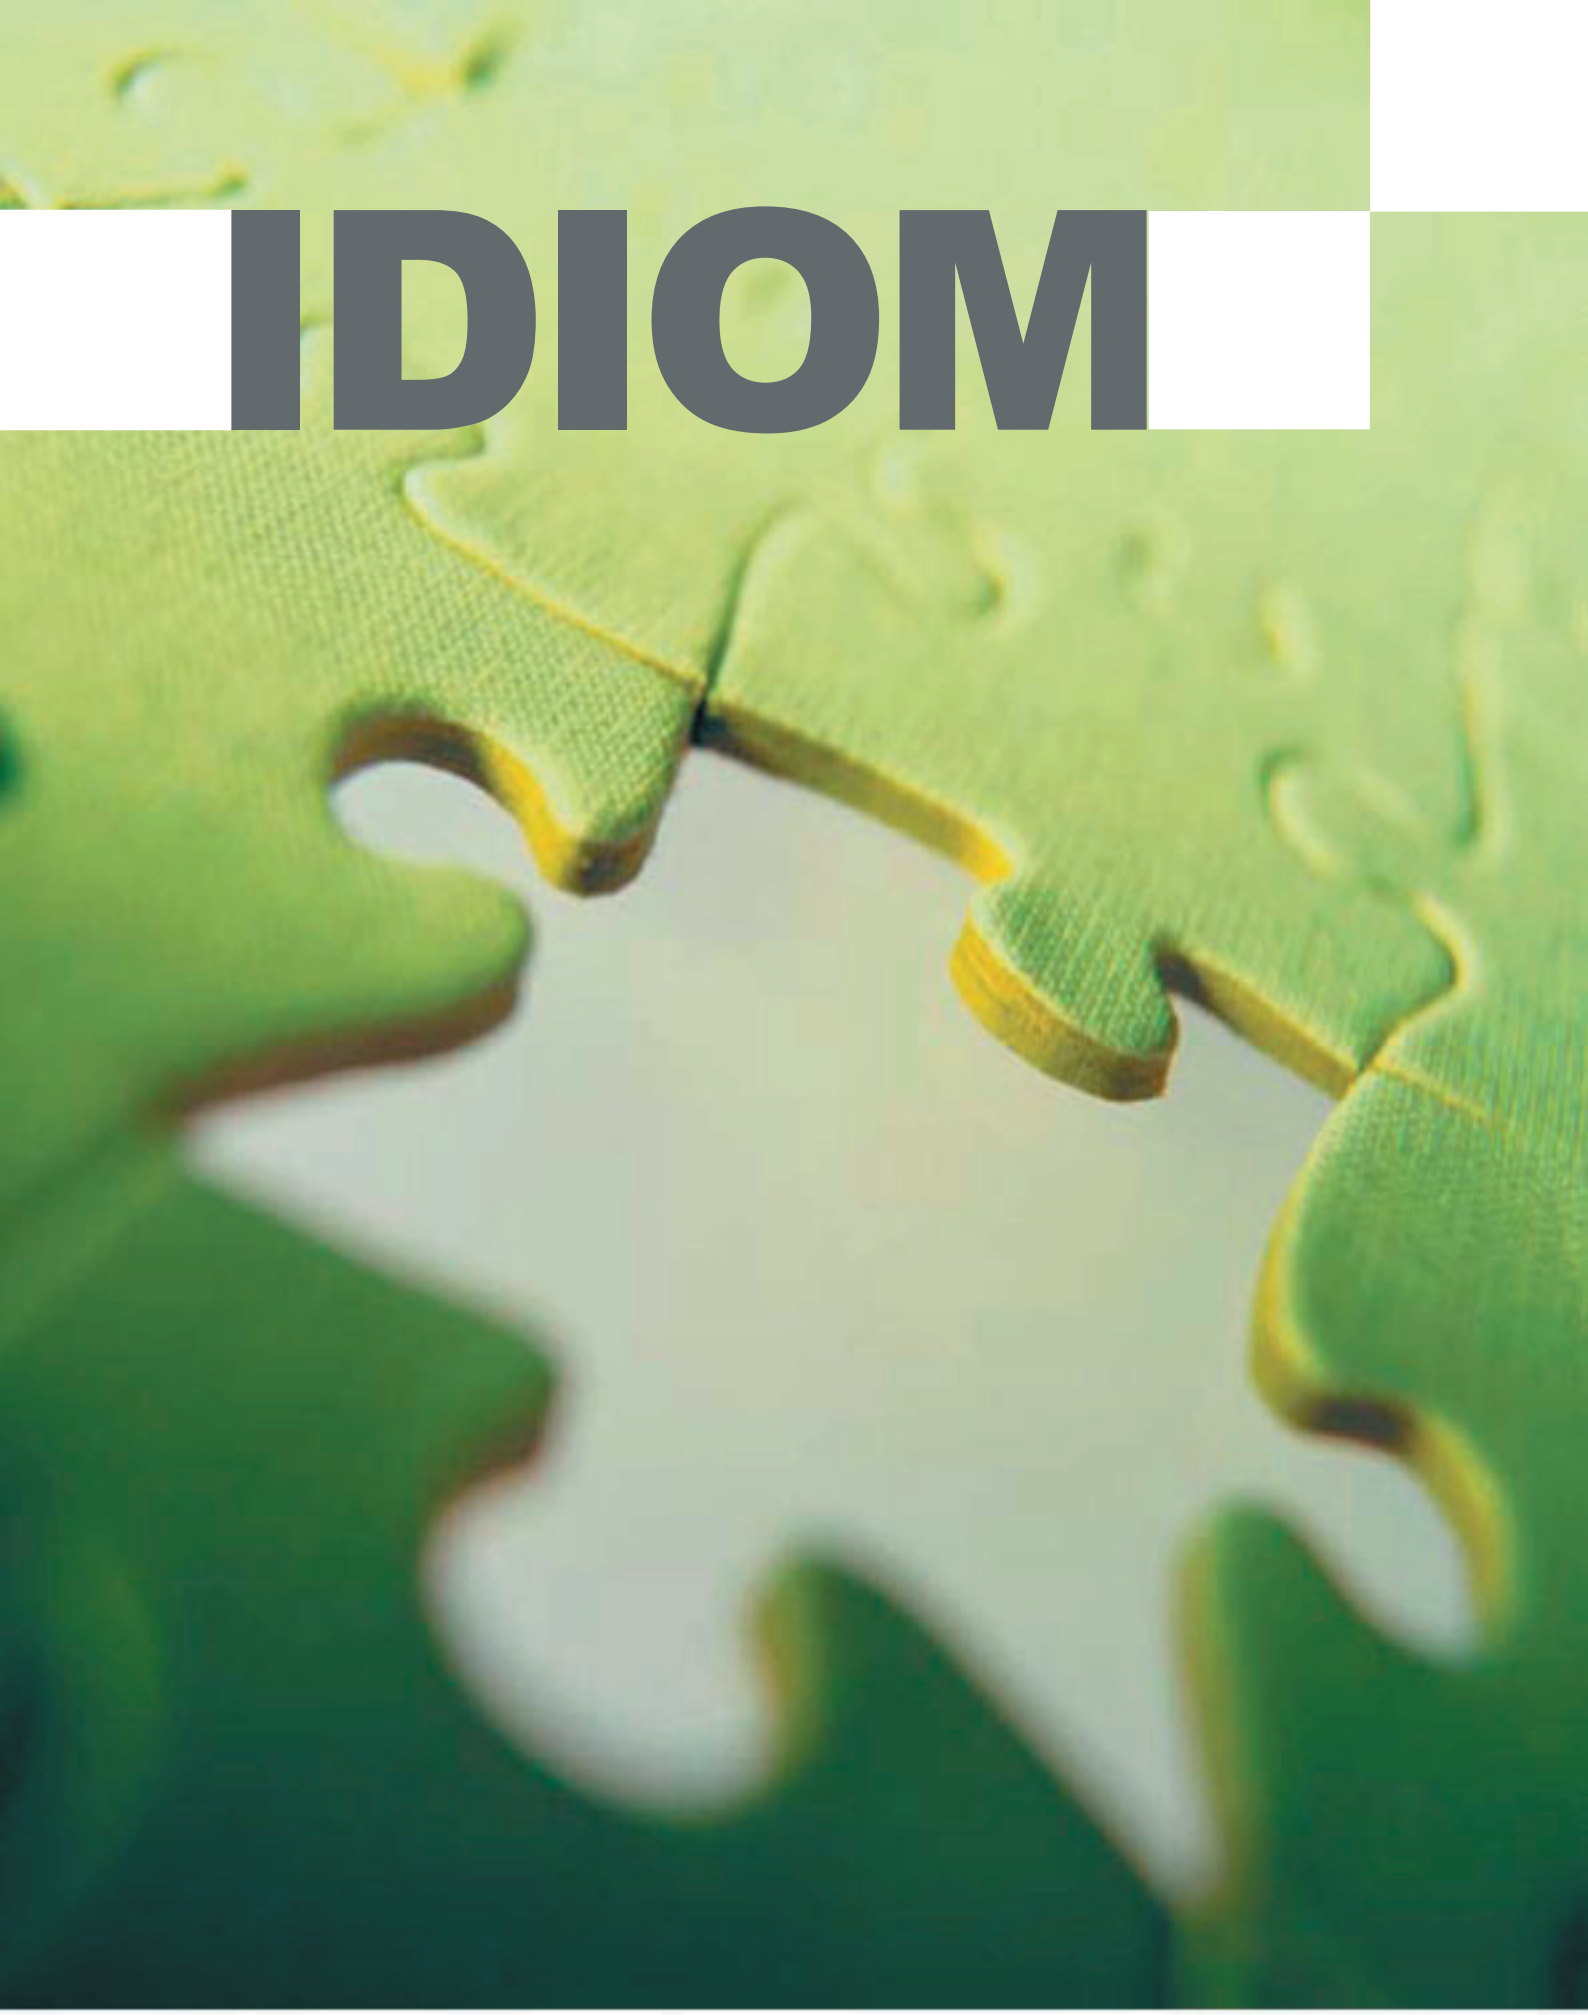

## Decision-centric Rules Modelling

## Introduction

The purpose of a business rule is to enable a business to consistently apply business policy to decision making in order to achieve desired outcomes. The decision-centric approach to business rules definition involves modelling the decision-making process as an explicit part of the business rules model.

The decision-centric business rules model comprises the following:

a. Fact (or data) model
b. Business rules, each of which determines a single value
c. Decision model
d. Versioning by date and business scenario

The fact model is a tree-like hierarchy of relationships between facts (represented by one or more XML schemas).

The decision model is a tree-like hierarchy of processing relationships between decisions (represented in the form of a mind-map). The decision model proceeds from the level of complex decision making that determines many facts, to the level of atomic decision making that determines single facts.

Atomic decisions are linked to both the fact model and to the business rules. Complex decisions are linked only to the fact model and are not directly connected to specific business rules because the outcome of a complex decision is the sum of the facts derived by the atomic decisions that comprise the complex decision.

The purpose of the decision model, as distinct from the specification of individual business rules, is to:

a. Enhance the business's knowledge of how it applies its business rules within the decision-making process.
b. Reduce the granularity of the decisions to a level where it is practical to define a specific rule that determines a single-valued outcome.
c. Provide a level of abstraction that enables reuse of rules across atomic decisions and enables versioning of the rules by time or by business context (such as `geographic location' or `customer').
d. Enhance the validity of the overall business rules model by reconciling the decision model with the fact model.
e. Provide a new software component, the `decision model', which extends business capability by adding intelligence and configurability to automated processes.
f. Provide a visual, business-accessible representation of the decision-making process that implements business policy.

IDIOM views a business rule as a calculation or formula that determines the outcome of a single atomic decision. A formula can be composed of many logical, mathematical and other operations, provided it determines a single- valued outcome.

## The Inference approach

The inference[1] approach to business rules modelling is drawn from knowledge engineering theory. A program that uses the inference approach to modelling business rules is known as an *expert system*.

An expert system contains knowledge derived from an expert in some domain. This knowledge is used to help individuals using the expert system solve some problem in that domain. An expert system whose knowledge is expressed in rule form is called a *rule-based system*.

An expert system has the following components:

[1] To clarify, we do not intend the term 'inference' to imply all possible algorithms that provide an answer. We specifically intend the term `inference' to exclude algorithms that are comprised entirely of formally declared logic.

a. Knowledge base – domain-specific problem solving knowledge
b. Facts – what we know at any time about the problem that we are working on
c. Rules – atomic relationships between the facts
d. Inference Engine – applies the knowledge in the knowledge base to problem solving

An expert system employs an inference engine to use information in the knowledge base to deduce a conclusion when given some facts. It is an algorithm (or series of steps) used to solve a problem. A knowledge base is a collection of information stored in a structured fashion.

The inference approach to business rules modelling involves atomic rules (rules that cannot be broken down any further) and an inference algorithm that determines, from known facts, which of the atomic rules should be applied, and in what order, to determine a new fact or facts. This algorithm does not need a human to directly express the steps required to determine those facts. The rules are in the form IF <some conditions> THEN <some consequences>.

Inference in a rule-based system involves:

- Selection of rule candidates that can be used to solve a particular problem
- Choice of one rule to execute at each step in the process of deducing a conclusion
- Deduction of a conclusion

Inference is achieved using a forward-chaining algorithm or a backward-chaining algorithm.

A backward-chaining algorithm works backward from a conclusion to be proved to determine if there are known facts that can be used to prove the truth of the conclusions.

A forward-chaining algorithm works from the initial facts to infer the conclusion (i.e. concluding fact). A commonly used forward-chaining algorithm is the RETE algorithm.

## RETE Algorithm

- Used in forward chaining rule-based systems.
- Performs very well in situations in which very few changes occur (i.e. where there are few new rules or changes to rules).
- Has significant storage overhead.
- The RETE algorithm and its derivatives are the only algorithms for implementing production (inference) systems whose performance is independent of the number of rules.

A hybrid system can combine both the backward-chaining and forward-chaining approaches to inference. Many rules-based systems used for business rules modelling are hybrid systems.

The purpose of the inference approach is to maximise the utility of the knowledge gained by the expert system. The utility of the knowledge is enhanced by using inference to allow rules to be applied to situations that were not necessarily foreseen at the time the system was developed. Actors that use the system can then use this knowledge to help solve problems in a particular domain, without the system having to be designed around the particular problems that are to be solved.

## Decision-centric vs. Inference approaches to rules modelling

A rules-based system suits applications that require inference. For other applications, a business may gain more knowledge by employing a system that directly models the algorithms for rule execution. The decision-centric approach directly encodes the algorithm for rule execution by modelling the decision-making process. This visual model is easy for a business user to understand.

## Independent modelling of the business algorithm and rules

The preferred algorithm employed by a business to execute its rules is not necessarily the same algorithm determined by an inference engine. Explicit modelling of an algorithm for rule execution means a business has gained explicit knowledge of the exact algorithm it uses. In addition, the algorithm can be modelled independently of the business rules. It then becomes possible for rules to be applied in an efficient manner i.e. one that does not require redundant re-application and which enables coherent versioning of the rules.

The knowledge gained and declared through the process of building the decision model may be as valuable as the knowledge that is exposed by the individual rules underlying the decisions. For instance, the decision model may expose gaps in the decision-making process that arise from a poor understanding of the business or from a lack of formal, coherent rules underlying the decisions. Furthermore, this knowledge serves as a basis for changing from a rules model that is incomplete, or oriented on a technical process, to a rules model that is complete and oriented to a business process.

It is important that the validity of the business rules model is maintained for as long as the policies that it implements continue to support the objectives of the business. Developing and enhancing validity is important because it results in fewer errors in the business rules and in more accurate decision making in terms of declared business policies. The validity of the business rules model is enhanced through the use of a decision-centric approach.

The decision-centric approach requires the business rules to be modelled in the context of both the facts that those rules require and the decision-making process that will utilise those rules to determine new facts. In contrast, the inference approach expressly ignores the decision-making process and its role in describing in a meaningful way whether, when, and under what circumstances, individual facts are to be determined.

There are situations in which the precise decision model is unknown. From a generic knowledge engineering perspective, it is desirable that an inference engine handles these situations. However, in many business contexts the business may prefer to refer the decisions in these situations to a human for resolution. The decision-centric approach can precisely determine the situations that require referral. As the referred decision making becomes better understood, so it too can be declared back into the decision model and thereby automated, so that the boundary between declared and inferred/referred decision making is always moving.

## Many applications do not require an inference engine

Many applications for a rule-based system do not require an inference engine. Consider the calculations involved in deciding how much money is owed in a tax return and who the money is owed to (either you or the government). The government has laid out a series of calculations to determine the amount owed and a rule based on the calculated amount to determine who the money is owed to. The algorithm for selecting the calculations or rules to apply to the tax return is well known, as is the order in which to apply the rules. Therefore, it makes sense to directly encode this algorithm. In a business information system, it is usually clear from the business context when a particular algorithm should be applied to the problem at hand. For instance, if an employee requires the price of a product, they can determine this using the product pricing system. The data will be keyed in and the pricing system will execute an algorithm that applies known rules in a well-known order to determine the price. An inference engine is therefore not required in this case. In our experience, many, even the majority of business contexts that involve information systems, do not require an inference engine.

The inference approach enhances the utility of the knowledge gained by the expert system. However, if the business already knows its decision-making process and the rules that underlie its decisions, this enhanced utility of knowledge is redundant. In addition, the enhanced utility of the knowledge

carries with it the risk of unforeseen and undesirable interaction between rules in the knowledge base. The decision-centric approach allows for broad flexibility in the business rules because the only absolute requirement is that a business rule will produce a single-valued outcome. This wide ranging scope is more important than it would seem at first appearance, because it allows a business domain expert to directly capture their own knowledge in a model form, rather than abstracting it via a series of facts and rules that are re-factored by a software rules engine.

As an important side effect, this very close, real-world relationship between the business domain expert and their decision model makes validation and testing of the model much easier for the business domain expert, and less risky and less costly for the business overall. This gives the owner of business policy the hands-on ability to tightly couple the policy with its computerized implementation.

## Summary

The inference approach:

- Considers that only the fact model and the atomic business rules are part of the knowledge to be gained in relation to its business rules model.
- Cannot expose knowledge as to how and why a business applies its rules in a particular business context.
- Uses the inference engine to handle uncertainty with regards to the rule execution algorithm, rather than business experts.
- Uses an inference engine that is redundant in most business contexts.
- Breaks calculations into many atomic rules that don't coherently describe the complete calculation from a business perspective.
- Allows reuse of atomic rules only.

The decision-centric approach:

- Allows the business to gain explicit knowledge of the decision-making process itself.
- Improves the business's knowledge of decisions that range from simple to very complex.
- Enhances the business rules model by grouping the decisions in ways that reflect the actual use of the rules in the business.
- Reduces the complexity of constructing the business rules model by refining the decision-making process to a point where a rule with a single outcome can be defined.
- Allows an `atomic' rule to be defined by the business `idiom' i.e. the rule is expressed in a language that is natural for the business.
- Captures and presents all the facts and logic required to make a decision in a given business context, resulting in a picture that is meaningful to the business.
- Enhances the overall validity of the business rules model by requiring reconciliation between the fact model, the decisions, and the rules.
- Allows versioning of the rules to be modelled explicitly and separately from the rules.
- Allows reuse of both atomic and complex rules.

In comparison with the inference approach, the decision-centric approach results in a greater gain for the business as the result of producing a more complete and valid business rules model that is aligned with the business `idiom'.