

Technology

Process

Compliance

» White Paper

Using Automated Decisioning and Business Rules to Improve Real-time Risk Management

Sandeep Gupta, Equifax
James Taylor, Smart (enough) Systems

August 2008

Table of Contents

- Introduction to Decision Management and BRMS.....1
- Inherent Risk Requires Sophisticated Decisioning.....1
- The Power of Decision Management.....2
- Automated Credit Risk Decisioning.....6
- Decision Management and Credit Risk.....6
- Optimizing Strategy Management.....9
- Transition from Technology to Business Strategy13
- Suggested Reading.....14
- Equifax14

Introduction to Decision Management and BRMS

Inherent Risk Requires Sophisticated Decisioning

Streamlining processes and improving decision making is critical as demands intensify to minimize credit risk, increase customer acquisition, and improve retention rates. In today's fluctuating economic environment, companies, especially financial institutions, stand to make strong gains from decision management if it is applied correctly. This requires the integration of robust, automated decision management systems into operational processes. However, these decision management systems must be grounded in the context of a sophisticated Business Rule Management System (BRMS) in order for companies to gain the agility, profitability and new growth to stay competitive and successful in the marketplace.

This white paper includes an introduction to decision management, written by James Taylor, one of the leading experts in the field. This introduction delves into predictive models, business rules and Champion Challenger testing as elements of successful decision management.

This paper examines the relevance of decision management to credit risk, how business rules can be defined and managed and how what-if analysis and Champion Challenger testing improve results. You will learn:

- What decision management is and what it requires
- Why decision management is essential for effective credit risk management
- Options for editing business rules
- Process and discipline for improving automated decisioning; and
- Approach for moving to business strategy management.

This white paper is written for credit risk managers, technologists involved with decision management systems, and others interested in understanding more about credit decisioning.

Equifax has helped the world's largest banks and telecommunications companies implement rules-based strategies that achieve an optimal balance between risk and new account acquisition.

The Power of Decision Management

Napoleon Bonaparte said, “Nothing is more difficult, and therefore more precious, than to be able to decide.” In today’s 24x7, customer-centric and hyper-competitive world, the time and resources available to make decisions are more constrained than ever. Operational decisions – those required to process high-volume transactions in front-line systems and processes – are under particular pressure. Decision management – automating and improving these high-volume operational decisions – is the only way organizations can ensure that every operational decision is taken correctly, profitably and legally. Decision management develops automated decision services using business rules to automate those decisions, adds analytic insight to these services using predictive analytics and allows for the ongoing improvements of decision-making through Champion Challenger testing.

To maximize the efficiency of a business process, you need as many automated decisions as possible. Correctly done, this automation eliminates subjectivity and enforces compliance with regulations such as Sarbanes-Oxley while increasing straight-through processing. Decision management boosts worker productivity, adds consistency to the way you conduct business, and manages risk more precisely.

There are four key elements to decision management. It’s about finding, automating and improving decisions to improve your business. It relies on business rules and business rules management systems to make it easier to see what is happening and easier for the business to control it. It uses data mining and predictive analytics to find the right rules and to make those rules “smarter” by inferring interesting things from the data you have, and it’s about ongoing improvement using Champion Challenger.

Explicit decisions

The first step in decision management is to find and make explicit the decisions that drive your business. Many of these decisions involve the assessment of risk for a particular transaction, customer or action. By making these risk decisions explicit, you can make them:

- Flexible for increased agility and currency
- Visible and measurable

- Business driven and owned
- Consistent across channels, systems and processes
- Compliant with regulations and policies
- Customer-centric

You can also ensure that these decisions embody your strategy and company values as well as the expertise and “tribal knowledge” of your experienced staff. Explicit decisions, as we will see, also give you the opportunity to truly leverage your corporate data assets.

A Business Rules Management System or BRMS is a complete software platform for defining, managing and executing business rules in an information systems environment.

Business Rules and BRMS

To make it possible to manage explicit decisions, you need a way to specify them clearly and consistently. Business rules are used to define the policies, regulations and best practices that should be used to make a decision. If the business rules for a decision are to be automated, and they will need to be for most operational risk decisions, then they must be implemented so they can be easily managed and edited. This requires the use of a Business Rules Management System or BRMS. Using a BRMS allows business rules to be clearly defined and correctly executed. Automated decisioning works like a well-oiled machine when the rules being executed are defined without any ambiguity.

Automating the implementation of business rules also makes policy-makers’ lives much easier. They can accurately measure whether or not their policies are working by eliminating subjective manual enforcement and, using the wide range of user interfaces provided by a BRMS, they can actually take control of the rules directly. This allows them to edit and refine them as necessary.

To achieve a high degree of automation in credit decisioning, risk managers need a BRMS as well as partners experienced in the definition of decision rules to improve automated decisions; proper tools to assist with the analysis, definition and execution of decisions; and appropriate data to analyze the results of applying their risk policy.

Data and Analytics

Business rules must work with data and the quality and availability of that data is critical. The broader the range of data that can be used, the wider the range of rules that can be

written and the better the quality of decision is likely to be. Data may consist of feeds from an external system (such as user applications for credit cards) as well as data extracted from sources that provide additional information (corporation lists, credit files, etc.). All this business information must be defined and managed so that it can be used to define business rules.

Many times the input data is further aggregated. Generally these aggregation formulas are also defined as business rules. An example of such a rule would be one that calculated total outstanding balances unpaid within 30 days of invoice for bank cards opened within 180 days. These rules define calculated attributes which are then available to subsequent rules.

Historical data also has a role to play. Using data mining techniques, you can analyze historical performance to define rules and thresholds that make sense for ongoing operations. Using established techniques to analyze historical applicant data to build a decision tree, for instance, can help segment new customers into groups that are likely to respond differently to different treatments. Decision trees can be easily represented as business rules, allowing them to be easily integrated into automated decisions.

For a more sophisticated solution, a much higher level statistical model exists, known as a score model or score card. Score cards can be built by individual organizations on their own data or by companies like Equifax against pooled data from multiple companies. VantageScore® is one such example of a pooled data score card.

The definition of these models involves more intense analysis of historical data. Data is analyzed to see which attributes of a customer or application contribute to the riskiness of that particular application. The extent to which each attribute value contributes risk is modeled and a score card created. These models can be represented as a set of business rules, but realistic implementation of these models in an automated decision requires the ability to specify the rules using a custom interface specifically designed for score cards. With score card rules defined in the same environment as all other business rules, you have a complete picture as well as support for an end-to-end revision process.

The precision of the resulting rules and models is directly proportional to the available data, the defined scenarios, and the capabilities of the tools being used. SAS and several other such tools are widely used to analyze sample populations and to develop predictive models and rules. When using data mining and analytics to improve the automation of decisions, it is best if you have information both about customers accepted by your organization and about those not accepted. For most organizations automating risk management, however, data about rejected applications is not readily available for analysis. The data may not have been collected at all or it may have been collected too inconsistently to use. In these circumstances, the use of pooled models such as VantageScore is recommended.

Champion Challenger

Champion Challenger testing is used to continually refine and improve decision making. Decisions require constant assessment and measurement for a number of reasons. First it can be hard to tell at decision time what is a good and what is a bad decision. For instance, bad credit risks may not show up for some time after their application is initially accepted. By the time it is clear that someone is a bad credit risk, it is too late to go back and see if a different decision would have resulted in lower losses, for instance. Secondly, it is often hard to analyze the potential impact of a proposed change for similar reasons. If the proposed change would treat some customers in a new way, there may well be no data about how those customers will respond to that treatment and no way of telling what impact this behavior will have on risk.

Champion Challenger addresses both of these issues. It allows for the definition of champion and challenger paths at any step of the decision. The “champion” is the one that will be applied to most transactions and is either the current approach (if a change is being considered) or the approach you expect to do best (if you are trying to gather information for later comparison). The “challengers” might be new approaches you wish to test or alternatives that will let you gather data for future analysis. Champion Challenger always uses only one champion path, but one or more challenger paths can be defined. Each path has unique decisioning logic – business rules, score cards and so on.

BRMS technology enables customers to develop “what-if” strategy testing of new scores and rule matrices using their own off-line and archived data.

Although both champion and challengers are deployed in production, the challenger strategies are generally executed for a smaller set of applications. Once a new strategy has been validated, then it can be applied to the entire production. Doing so ensures that if the strategy update does not produce the desired result, then the impact is limited to a much smaller population. Similarly, if a champion strategy is not getting the results desired, challengers that have been running in parallel will offer clues as to what might work better.

Automated Credit Risk Decisioning

Decision Management and Credit Risk

Credit risk decisioning in the BRMS world integrates relevant credit and non-credit information from multiple data points, including scores, ratios, models, fraud tools, judgmental criteria and policy compliance rules, and allows for real-time credit decisions. This empowers the front-line employees at banks and retail stores to offer real-time credit-granting decisions at the point of sale.

Real-time risk decisioning has become increasingly complex, with the need for actionable customer intelligence delivered from both internal and external databases. A good BRMS offers the following:

- Centralized and automated risk decisioning processes across all sales channels
- Consistent business decisions that are regulatory compliant across the enterprise
- The ability to quickly modify centralized business rules or credit policies to accommodate changing market conditions.

Many legacy systems use hardwired business rules, which makes it extremely difficult to deploy centralized business rules across the organization when circumstances change – new business acquisitions for example. IT departments can’t easily manage this change process. One solution is to deploy hosted or ASP BRMS systems which integrate online with legacy systems. The business rules for today’s automated credit risk decisioning system now reside outside the legacy system.

Editing Business Rules for Credit Decisioning

As noted above, business rules are critical to effective decision automation and management. Because business rules represent the point of intersection between business know-how, analytic insight and systems development, it is crucial that they be accessible to multiple audiences. Business and technical users must be able to represent business rules using several different metaphors, such as “if-then-else” statements, decision nodes in a rules flow diagram, decision tables, cross-tabs, and decision trees. A good rules editor must have a provision to define these business rules using any of these metaphors.

The first metaphor is that of rules flow. A rules flow allows the definition of the steps within a decision that must be executed to make the decision correctly. Each step might consist of a set of rules or a score card to execute. By allowing steps and branches to be defined, the BRMS can handle all aspects of a decision, improving agility and completeness.

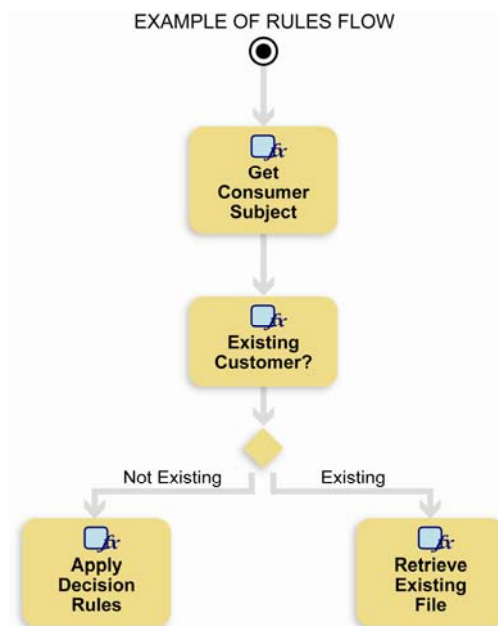


Figure 1: Example of Rules Flow

It is also fairly common to have multiple task points defined as part of the rules flow. For example, if business rules reach certain conclusions, then they should perform a specific task such as access a defined data source, mail a letter or send an e-mail. Although the performance of these tasks is not the responsibility of the business rules tools, they must provide a mechanism to depict how these activities will take place.

Having defined the overall flow, business rules will need to be defined for the various steps within the flow. Business rules might be defined as If...Then...Else conditional statements, each independently managed by the BRMS, or as a table of rules, sometimes called a decision table. Other options include template-based rules and decision trees.

TYPES OF RULES

Conditional Rules:

Duplicate Sequence #*

[Definitions]

If
 Duplicate Sequence #? : true
 and Account is between : 0 [±] and : 4 [±] days old

↳

Then
 set transaction's decision state to : Decline -Duplicate App

↳

Table Rules:

FICO Score		Decision	Reason Code (work list)
min	max		
0	450	Policy Decline	D100 - Low Score
451	620	Manual Review	M208 - Review w/Credit Mgr
621	700	Approve - Basic Offer	A314 - Request Copy of DL
701	760	Approve \$2000 Limit	A303 - Offer Add'l Svcs
761	800	Approve \$2500 Limit	A303 - Offer Add'l Svcs
801	999	Approve \$5000	A303 - Offer Add'l Svcs

Figure 2: Types of Rules

Once the initial set of business rules has been defined using a tool and deployed in production, a mechanism must exist for revisions. Ease of revision and the resulting business agility is a primary benefit of using business rules and a BRMS.

Many times these changes are obvious and do not require extensive validation or testing prior to deployment to the live production system. This is especially true if the BRMS has been used to control the editing of these rules and if it provides a mechanism to promote such changes to production without any IT involvement directly by the policy managers.

In other cases, a change will require a degree of validation and testing prior to production. By considering how the updates come about – the analysis of input application data and

results, for instance – you can determine the different degrees of validation and testing required.

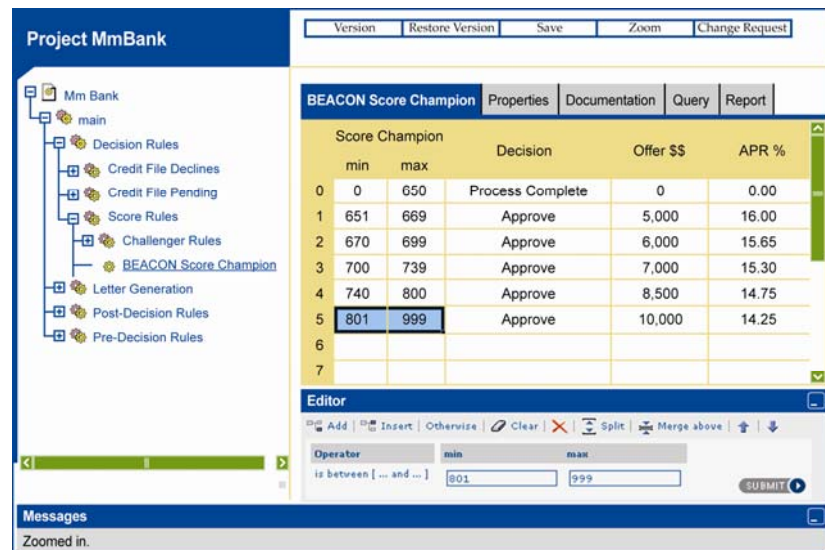


Figure 3: Business User Editing of Rules

Optimizing Strategy Management

The flexibility and ease of management provided by business rules and a BRMS can and should be exploited to continuously optimize business strategy. Champion Challenger testing is one approach while off-line hypothesis testing is another.

One of the most effective ways to optimize strategies is to constantly challenge the status quo. Lenders who become complacent and allow the same strategy to be used for extended periods are likely to be outperformed by more aggressive competitors. Champion Challenger testing is ideal for those who want to continually refine their strategies. Champion Challenger combines real results - actual customers are treated by the challenger strategies - with minimal risk.

With Champion Challenger testing, a planned change or an alternative being considered is deployed in production but applied to only a small set of applications. This enables you to gather information about the performance of the alternative strategy while limiting impact.

Rule 1 = current strategy:			Rule 2 = new strategy:		
Beacon Score		Decision	Beacon Score		Decision
min	max		min	max	
0	450	Policy Decline	0	450	Policy Decline
451	620	Manual Review	451	620	Manual Review
621	700	Approve – Basic Offer	621	700	Approve \$1000 Limit
701	760	Approve \$2500 Limit	701	760	Approve \$2000 Limit
761	800	Approve \$5000 Limit	761	800	Approve \$2500 Limit
801	850	Approve \$6000 Limit	801	850	Approve \$5000 Limit

Champion

Challenger

Figure 4: Champion Challenger Rules

Champion Challenger testing should allow for the implementation of gating criteria, allowing application processing flow to track which path was actually taken. An interesting problem exists where multiple challenger paths may be required to validate multiple updates to a single overall business rules set. One should be careful to ensure this type of implementation does not give rise to several possible paths of execution, as each gate will split the path. A more appropriate strategy, used prior to the next gate in a rules set, would require a split of the previous challenger to converge again while the same previous gating criteria would be used with consistency.



Figure 5: Setting gating percentages for Champion Challenger strategy

Champion Challenger testing allows for testing to be done in the live system. Another approach involves conducting predictive analysis with the sample data to arrive at a desired result. Prior to the application of any new strategy within the production system, even one being deployed as a challenger, an off-line analysis against the historical applications should be conducted. You can compare with the historical results to ensure that the change does indeed improve results.

This kind of what-if analysis allows for the testing of new scores, rule matrices and other hypothetical scenarios against off-line, archived data without impacting your production environment. You can effectively test the impact of any changes to score cut, policy changes, cross-sell assignment, data source assignment, decision assignment or line of credit assignment.

Here is a simple illustration of how rules editing, Champion Challenger strategies and hypothesis testing work together:

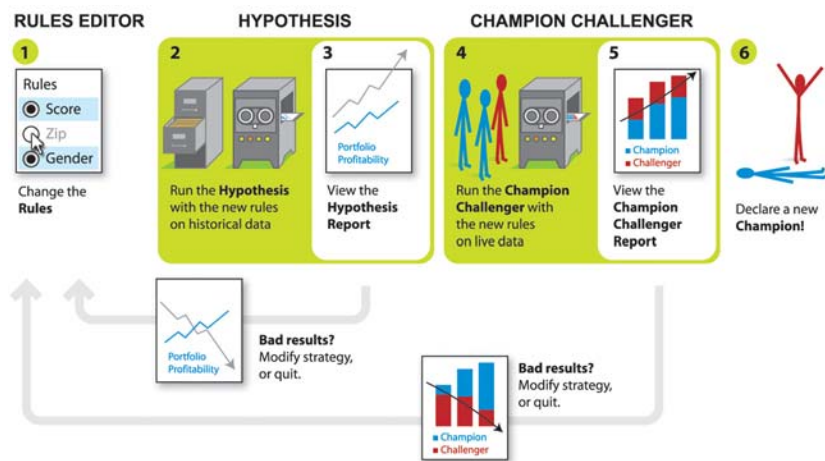


Figure 6: How what-if analysis and Champion Challenger work together

There are certain parts of a production rules flow that must not execute in the off-line environment. In each business case these elements will need to be replaced with “mock” services or simulators. For example, a particular production business rule may cause the generation of an e-mail. However, in the case of what-if analysis, an actual e-mail must not be created. Another scenario might be that a real credit application may be accessed. However, with off-line analysis, the previously used credit file must be reused.

Re-use challenges exist when accessing archived data because the actual information may be dated and must be adjusted so that when the rules are executed, the comparison is done in an appropriate context. A good example of this involves an activity in a consumer's credit file that may be analyzed to determine its age. When this data is analyzed in an off-line environment six months later, it may produce different results. For this reason, either the data in the credit file must be adjusted for the age of the data, or the context of the execution must be adjusted to move the time back when the actual execution had taken place in production.

So, for an effective off-line or what-if analysis for the validation of an enhanced predictive model with historical data, an appropriate tool must facilitate proper execution as defined above. Not only that, the facility must also assist with the creation of alternate business rules scenarios and with the identification of data sets from the archive to test against the entire population. Testing against a target sample may be hindered with the analysis of the rules change's effectiveness.

The biggest question still remains around how to develop suggested changes to the business rules. As previously stated, update suggestions generally are derivatives of historical data analysis combined with the feedback data (performance data) in order to compare how the decisions fared with the actual results. Determining whether a decision to approve a consumer for a credit card was a good result or a bad one is a classic illustration of this point.

At the basic level, a data mining tool may appear appropriate to conduct this type of analysis in order to review various characteristics of input data against the historical decision rendered as well as the performance results. Once analysis such as this is conducted, it may be almost impossible to define a strategy that eliminates all risks. Thus, the analysis turns to goal-based optimization. For example, if a wireless operator has a customer base of 50 million with a write-off of \$40 million, it may set a corporate target of reducing its write-off to \$25 million but not reducing the customer base below \$48 million. With this type of decision, risk is still present that either one or both targets may still be missed. Nevertheless, the objective now becomes setting the credit policy to achieve these targets.

This type of credit policy analysis and definition now becomes an exercise for operations research and more specifically turns into constraint-based optimization. In certain cases, it has been shown that even with a sample data set of 50,000, more than 20 million scenarios have to be evaluated in order to come up with a new business rules set to meet desired goals.

Transition from Technology to Business Strategy

In order to navigate the complexities of today's business landscape where economic, regulatory and compliance pressures continue to mount, companies are seeking a more robust decision management strategy to effectively and efficiently address their changing needs when it comes to streamlining the application and risk-decisioning process, as well as supporting ongoing change control and strategy development.

When implementing a decision management solution, companies should look for a technology solution that provides the following benefits:

- Easily edit business rules to respond rapidly to changing market conditions, without the need for programming support;
- Test and optimize new business rules before rolling them into production.

With these capabilities, companies can focus on business strategy, not technical details. If a new business strategy can be implemented quickly and accurately, then business owners can continuously improve and refine their strategy.

In addition, a BRMS can go beyond just automating business decisions. For instance, a problem that has haunted telecom operators as well as financial institutions is how to optimize offers to their customers and eliminate subjectivity associated with Customer Service Rep (CSR) decisions. A decision made by a CSR regarding a best offer is likely not consistent across the entire company, because it is based on the subjective knowledge of the individual CSR. A solution to this problem involves using a BRMS to make consistent offer decisions based on the established business rules around such offers.

A national telecommunications company achieved comprehensive policy automation across all sales channels by working with Equifax to implement 150 multi-dimensional decision matrices, analyzing data from seven disparate sources.

These offers can further be optimized by leveraging online demographic and lifestyle data that helps customize even automated decisions around offers for new products and services.

Suggested Reading

“Smart (Enough) Systems: How to Deliver Competitive Advantage by Automating Hidden Decisions,” by James Taylor and Neil Raden. Copyright © 2007 by Prentice Hall.

“The Business Rule Revolution, Running Business the Right Way,” by Barbara von Halle and Larry Goldberg with guest contribution by John Zachman. Copyright © 2006 by Happy About®.

Case study: “Business Rule Management Systems and Financial Services: Equifax as a Case Study,” by Ariana-Michele Moore. Copyright © 2005 by Celent.

Equifax

Equifax has proven experience helping customers implement best-of-breed automated risk decisioning processes. Equifax’s InterConnect® platform offers specialized state-of-the-art business rules management for credit risk decisioning. Equifax leverages the combined value of our predictive sciences team, along with our professional services team, to help customers better evaluate their historical loan decisions and credit policies to increase automated decision rates and improve credit risk management. Please contact your Equifax sales representative to find out how Equifax can help your business implement real-time risk decisioning.

About the Authors

Sandeep Gupta is Vice President, Strategic Software Development for Equifax. Mr. Gupta is responsible for the management of software development in the B2B space for the InterConnect suite of software products at Equifax.

Mr. Gupta oversees the development of tools to provide solutions in identity management and risk-based application processing, including the use of business rules technologies to

automate decisions and improve maintenance of decision rules.

James Taylor is one of the leading experts in decision management. Author of the first book on the subject, he is Co-founder of Smart (enough) Systems, a consulting company dedicated to helping companies adopt and exploit decision management. You can reach him at james@smartenoughsystems.com You can also find out more about decision management at the Smart (enough) Systems blog: www.smartenoughsystems.com/wp.

Equifax, VantageScore and InterConnect are registered trademarks of Equifax Inc. All other trademarks and service marks not owned by Equifax Inc. or its subsidiaries that appear in this publication are the property of their respective owners. Copyright © 2008, Equifax Inc., Atlanta, Georgia. All rights reserved. Printed in the U.S.A.

EFS-797-ADV—8/08