

Next Generation: Using SAS® Decision Manager to Modernize and Improve Your Operational Decisions

Charlotte Crain and Steve Sparano, SAS Institute Inc.

ABSTRACT

This paper takes you through the steps for ways to modernize your analytical business processes using SAS® Decision Manager, a centrally managed, easy-to-use interface designed for business users. See how you can manage your data, business rules, and models, and then combine those components to test and deploy as flexible decisions options within your business processes. Business rules, which usually exist today in SAS® code, Java code, SQL scripts, or other types of scripts, can be managed as corporate assets separate from the business process. This will add flexibility and speed for making decisions as policies, customer base, market conditions, or other business requirements change. Your business can adapt quickly and still be compliant with regulatory requirements and support overall process governance and risk. This paper shows how to use SAS Decision Manager to build business rules using a variety of methods including analytical methods and straightforward explicit methods. In addition, we demonstrate how to manage or monitor your operational analytical models by using automation to refresh your models as data changes over time. Then we show how to combine your data, business rules, and analytical models together in a decision flow, test it, and deploy in batch or real time to embed decision results directly into your business applications or processes at the point of decision.

INTRODUCTION

The next-generation business infrastructures in operation today must accommodate machine-generated data from devices, sensors, and social media; inexpensive storage; changing business requirements, such as new regulations, new policies, or competitive changes; and the need for accountability for how results are obtained. These de-facto realities of 2015 have dramatically altered the way business is required to operate today. In addition, “analytics-based processes are commonplace in organizations, but all too often are ad hoc or stand-alone applications.”⁴ Businesses must respond to their customers regardless of external influences of market demand, rapidly changing regulations, and policies. Businesses require increased speed from development and test to production and results; more robust, predictive models and scoring on larger data sets or tables; model governance; self-service capabilities; and a repeatable process that supports decision making. See how to use SAS Decision Manager to modernize, improve, or facilitate operationalizing your analytics.

RIGHT-TIME DECISIONS

Increasingly, IT wants to give control back to the business users for increased and flexible development that meets the needs of the business. IT wants to collaborate with business users for deployment and support for batch and near real-time. Yet, as business users gain more control over business applications, it continues to be difficult to extend analytics to operational processes. ¹ Historically, the development of analytical business processes is based on the skills and expert knowledge of employees, rather than on a formal business process. These talented people develop their own methodology around an existing way of doing business and they make it work for many years and in many places. Even today, many organizations, public and private, use this approach. The analyst produces models and IT deploys or operationalizes the models. The deployment process can have substantial delays due to code rewrites in Java or due to other technical processes required by the production environment. Finally, management incorporates the model results with other experience into the final operational decision making process. But the decision making process can lack support for robust, analytical models because the analyst is disconnected from the model that is running in production. There is no automated notification process that enables the managing and recalibrating of models, as needed. “Operational decisions relate to the organization’s internal functions and the decision management framework itself, such as when to retire an analytical model.”⁴ “The roles and activity phases of the business process are not ... [ideal]. Numerous handoffs, overwhelmed and understaffed teams, and rapidly changing regulations and policies all contribute to a less than optimal decisioning process.”³ These constraints lead to business processes that

are time-consuming and inefficient at using scarce resources. These existing business processes live on, despite technology changes, market changes, retirement of personnel, and expertise that is not properly transferred or documented.

However, the companies that have business processes that are improved or streamlined for analytics can leverage analytics for everyday business activities and create a more predictable and valuable organization that delivers or enables downstream consumption of results at the time of decision. This availability of results at the time of decision making is known as pervasive analytics and the process by which it is accomplished drives a greater need for collaboration within an organization. “Companies are ... constrained by a lack of integrated processes—operating within the department/functional unit level—which is a death knell to effective analytics across the organization.”¹

SAS Decision Management enables an organization to control, manage, and automate the repeatable decisions that are at the heart of its business by effectively applying business rules, analytics, and optimization technology with an approach for automating and improving high-volume operational decisions. These decisions can be comprised of data preparation + prefix business rules + data transformations + model formulas or algorithms + optimization of some kind + post business rules + integration to a production environment. In order to streamline or modernize an overall analytic business process, a unified platform that gives a consistent user interface for the business analyst with flexible, unified processes to deploy decisions, is required. Companies and organizations in the public or private sectors can modernize the analytical business processes that address the problems being discussed in this paper and can take advantage of the benefits of SAS Decision Manager.

MODELS AND BUSINESS RULES

Models do not make decisions; rather they provide valuable guidance (for example, probabilities and forecasts) that can be applied, along with other data and business policies, to the decision-making process. Using models, data and business policies together, businesses can create more intelligent actions related to various tasks, including underwriting loans, administering credit, managing risk, interacting with customers, setting product pricing, or even identifying and investigating fraud. The combination of model output data, other data, and business rules allows for action to be taken based on the guidance provided by the model rather than based on experience alone, without any analytical guidance.

The need to link models and rules quickly and accurately is recognized as valuable within businesses across industries. However, model and business-rule management are typically separate and disjoint processes with little collaboration between the teams performing these activities. In a rapidly changing, customer-driven world, a model that is 6 months old might not provide the desired value. It might cause lost revenue, poor communication with customers, missed market opportunities, or lack of compliance with increasing regulatory requirements. Compliance breaches often result in fines or penalties.

Being able to effectively and efficiently use individual expertise, as well as the power of analytical models within your current business processes, means the precision of models can be applied to your decision-making process. But many companies are not satisfied with the modeling techniques they use today, and “many ... companies have developed or are now developing new or additional analytical models that provide valuable insight into their customers, their products, and their markets. Unfortunately, many of the decision systems cannot maximize the power of analytics in the business processes at the point where the decisions are made.”³ This increasing number of analytic models and their use in critical decision-making processes require that the performance of these models be monitored to ensure that the most effective model is being applied at the right time.

MODEL MONITORING

A production modeling deployment environment should include a way to monitor how the models are behaving. Models are built with certain assumptions that require regular monitoring to determine if those assumptions are still valid. Has the data profile changed in a fundamental way? Is there a shift in the distribution of values that the models assumed were true but are no longer true? Due to the growing number of models, businesses and their potentially limited resources cannot continue to rely on manual methods to monitor the models that are driving their decisions. Automation and intelligent threshold monitoring can help.

Model performance thresholds should be defined across model projects, and when a particular threshold has been met or exceeded, messages can be generated and the models can be updated in an automated way. Businesses can continually improve personalized customer service in the form of relevant ads and promotions or create meaningful customer segments based on gradually built consumer profiles and in-store behavior. For example, a patient profile can be enhanced by knowing which patients might have a readmission due to their particular medical history, current medical condition, related conditions, and proximity to health-care providers.

After a data scientist or advanced analytics modeler performs an analysis, designs, builds and completes testing on a model that is ultimately deployed to production without code rewrites, the process must be repeatable. An established and streamlined business process for development and deployment can help ensure that each time a model is put into production, the business can have confidence in the process, and can trust the results to be available at the point of decision. Regardless of new or different data or any other changing conditions, the development to deployment process can be adapted and repeated. "Once tested and validated, the decision management process can be approved for IT to run in a production environment. 'It might run in a batch environment, a database environment, a web environment, or a service-oriented architecture.' The same process must be able to be deployed in many different places to be sure we are giving effective and consistent treatment across all possible channels."⁴

MODEL GOVERNANCE

Lastly and very importantly, having a process whereby there is some level of governance around model development, validation, testing, and publishing and data governance on which the model is built for auditing is considered a best practice. Financial, energy, governmental, and health-care industries require auditing and need governance. Governance also requires a repeatable process. Using a workflow for authoring and publishing rules and creating and publishing models provides a framework for governance within SAS Decision Manager.

MODERNIZATION

Modernization with SAS is defined as the continuous improvement of a customer's analytic capabilities using SAS. The value of modernizing SAS business processes and analytical deployments can be realized by using the best available business analytics software and supporting architectures for better accuracy, scalability, and governance. Reduced time for operational analytics deployment supports current state of a customer or market which leads to faster and more relevant offerings or more personal service and satisfaction.

Making faster decisions implies that the technology and the process gives decision makers easy access to results at the point of decision. This decision process enables speed, accuracy, and agility. Agility is the ability to add, update, and retire models or business rules when needed, versus using gut instinct or relying too heavily on other manual processes, checks, and rework for deployment.

MODERNIZING ANALYTICS FOR SCALE

The following example outlines one customer's struggle to streamline and derive value from an end-to-end process where SAS models are built and then deployed to production. The challenges arise when the company reviews the existing steps involved in getting their models into a business process so that they can use them in decision making. The goal is to reduce the time between recognizing an opportunity or problem and implementing models and rules in order to take an action to address the problem or capitalize on the opportunity.

The pace of change is accelerating in the following areas:

- Data velocity and volumes are increasing. The number of models and the speed at which they are produced are increasing.
- Model types such as machine learning and iterative algorithms are becoming more pervasive. The requirement to easily integrate with operational business systems is becoming more prevalent.

The changes related to one or two areas would be enough for an organization to deal with, but with all of these dynamic forces at work together, modern approaches are needed. These approaches introduce automation, intelligent monitoring, the ability to quickly react to changes in the market, and the ability to enact changes in the operational systems. The decisioning process comprises various stages, all of which require streamlining and modernization.

Data and modeling requirements, along with related changes for adaptation, are accelerating. Decision-making capabilities are required to scale as well, in order to keep pace, deliver value, and continue to enable businesses to seize relevant opportunities.

Continuing with our example, during the modeling stage of the overall process, many analysts spend around 80% of their time creating and managing data and about 20% of their time modeling. Even though data management and governance for enterprise data are crucial to building reliable and trusted models, the process is often still a work in progress. Data management for analytics provides approaches that address the diversity, quality, and availability of data in the appropriate form to the decision-management processes for model and business-rule development.

After models are built and tested, they must be deployed. The overall business process requires that the IT organization deploy the SAS models, and often the deployment personnel consists of Java programmers. Typically, the request from the business to the IT department documents the deployment and model requirements, and then IT recodes the completed SAS model in Java and moves through a testing and validation phase. This process can take up to 6 months or longer depending on the model complexity and on the IT resource skills and availability. In addition, the complexity of the models might be constrained. For example, neural network models are difficult or impossible to rewrite in Java. Any new or updated SAS model goes through the same IT process, before it is put into production.

Business rules that are associated with the model are written in Java or executed from an operational business-rules management environment. The business rules are not necessarily well connected throughout the business process or well connected to the models. The customer example presented in this section has described a manually intensive and disjointed analytical process, which many businesses still use today.

GETTING STARTED WITH MODERNIZATION

You can begin a modernization effort at various points throughout this process. SAS customers are well positioned to start a business process modernization effort because the SAS framework and foundation is the basis for all SAS applications. But any organization can benefit from modernizing its decisioning and analytical processes. The process can start at various stages. For example, a SAS customer might desire to modernize by starting with re-architecting the business process that addresses SAS code or SAS® Enterprise Miner™ model development or the SAS code that is used as the basis of the organization's business rules. Another starting point could be to streamline the business process for managing the models or the business rule logic that is often being changed via programming. This is a possible first step in the modernization process even though the models might not be SAS models and the rules might be in SQL.

Starting the modernization process at these stages enables you to tackle the problems of big data management, model management, business rule management, and decision management. Process improvements at these stages deliver more flexible and reusable business rules, rapid model deployment, and deployment of decisions to operational business processes.

By examining the need for modernization across decision management processes, it is clear that the process improvements can deliver value by making the following improvements:

- leverage big data for better decisions
- increase collaboration between business and IT
- enable discovery through analysis
- integrate models and business rules within decisions
- operationalize decisions rapidly

Leverage big data or as much relevant data as possible for better decisions: SAS® Data Management products deliver capabilities to cleanse, profile, and load data to make the data easily accessible to the decision management teams. Increasingly important is ensuring access to data in large, distributed data file systems so that it can be used for analysis in support of model and business-rule building and testing. SAS Data Management for analytics and business rules is complementary to the back-end extract, transform, and load processes that are already in place.

Increase collaboration between IT and business: Process efficiencies can be realized by reducing the number of manual steps and the number of handoffs. Apply automation to reduce lag time in the process and reduce the number of manual steps. Business users want to be empowered through model sharing, data access, and access to tools that increase their productivity by reducing reliance on IT for programming skills. Tools that allow the business teams to model business rules and simultaneously access the analytical models reduce the need for specialized skills.

Enable discovery through analysis: “Gut instinct or feel” and “rule-of-thumb” heuristics might be adequate when you do not have access to analytically sound techniques that facilitate model and business-rule creation. Analytical model creation is already prevalent but the speed at which models are required to be created is increasing and the number of models that can be created using factory-like approaches is also rapidly increasing. But business rule development can also benefit from this analytical approach. Techniques can be applied to analyze your existing business data to unlock the rules embedded in the data. By looking at target variables, techniques such as decision trees can discover business rules and deliver the rules to the business user to evaluate, adjust, and experiment with. These business rules can be combined with model output for more precise results.

Integrate models and business rules within decisions: How often have you developed code with one set of assumptions only to begin testing and discover that the outputs are not what was expected? Having that insight earlier would have resulted in better outcomes. Modelers often work in isolation from business rule analysts because they have been treated as different disciplines and placed in different parts of the organization. But when we consider that a decision uses both models and business rules, the decision modelers benefit from knowing the inputs and outputs of the models while developing business rules. With knowledge of the model’s results, the business rule designer can design and test business rules with model data, thereby creating decisions based on models and rules.

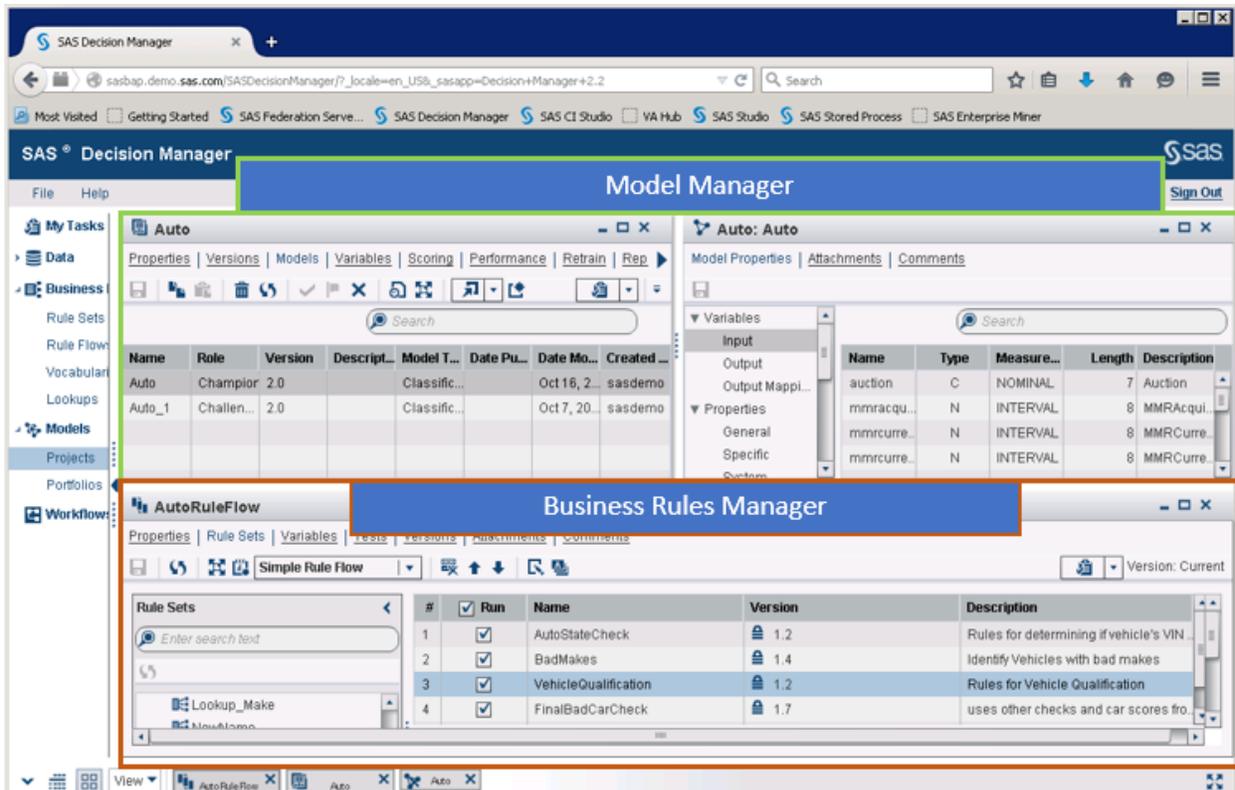
Operationalize decisions rapidly: Streamlined development and testing processes are important to modernizing the decision management process. Decisions must be implemented quickly because the value of the models and business rules can rapidly decay as data changes and customer needs shift. The ability to quickly deploy the decisions requires integration with the execution platforms, which can be varied. Execution platforms include the large distributed file systems, databases, and the web services that offer decisions as services. As data volumes and data stores grow, the ability to move the data to the decisioning engine becomes more difficult and costly. Deployment techniques must address moving the decision execution to where the data resides to eliminate data movement and deliver better performance. Speed also comes in the form of flexibility to implement new business rules that reflect regulatory changes, customer needs, and changing business policies. Responding quickly to these changes results in more adaptive and less brittle systems that better match desired actions in a changing market.

Next, we see how these various modernization needs are addressed using the SAS Decision Manager.

SAS DECISION MANAGER OVERVIEW

MODULES CONTAINED IN SAS DECISION MANAGER

Display 1 shows the SAS® Model Manager and SAS® Business Rules Manager modules in SAS Decision Manager that are used to create components for deployment in the decision process.



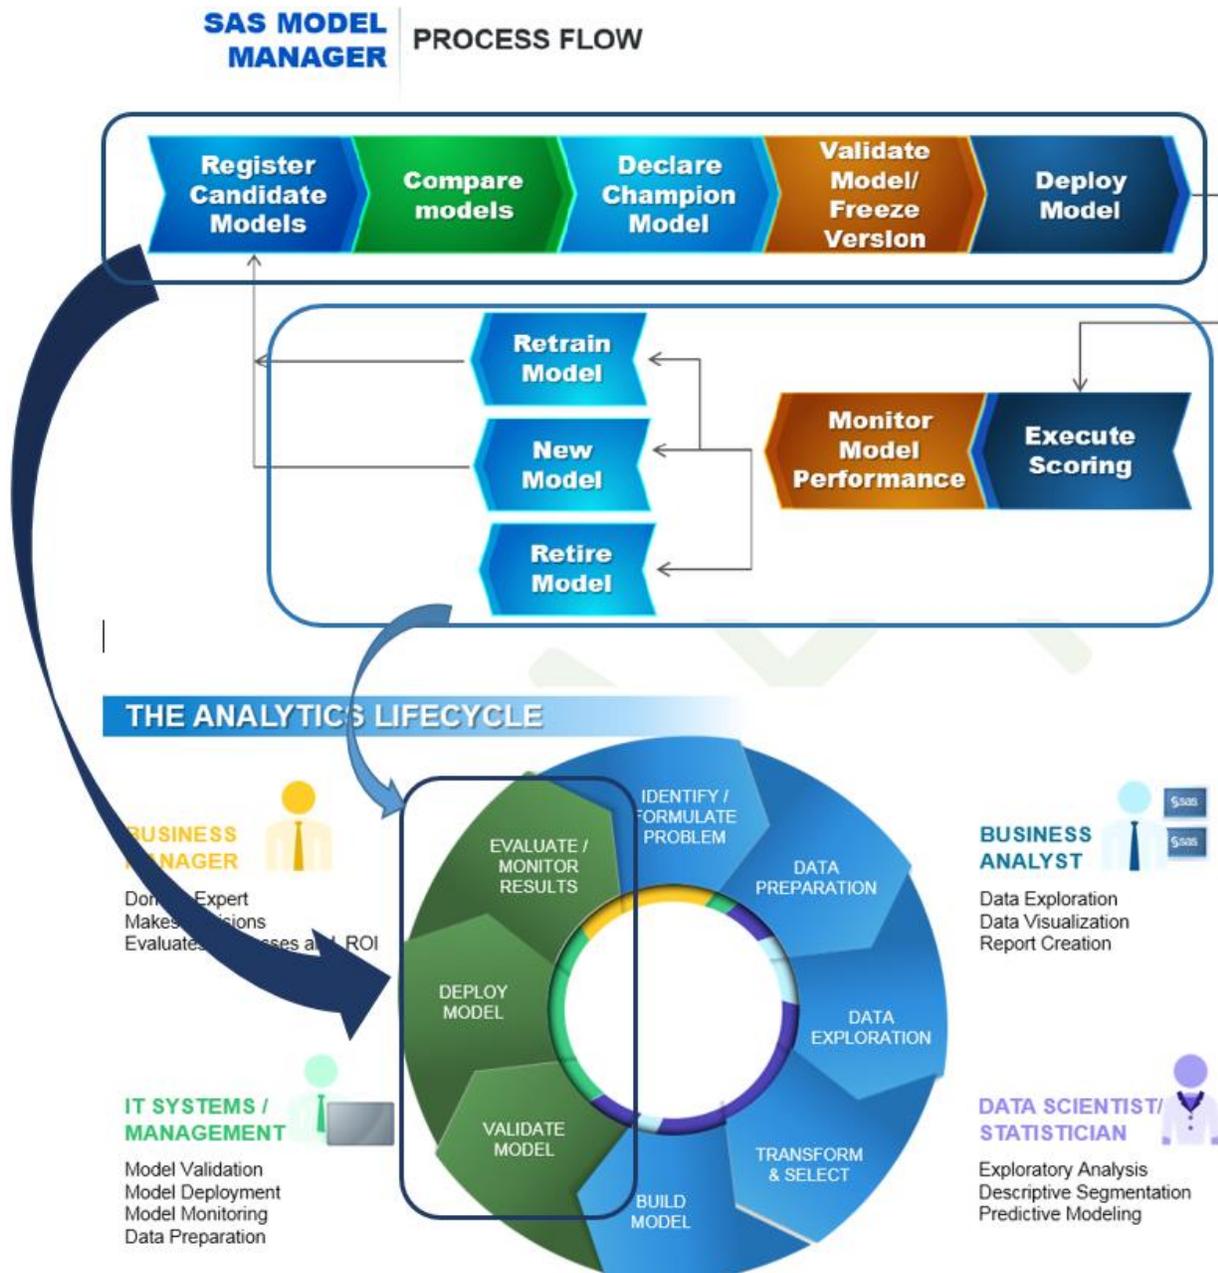
Display 1. SAS Decision Manager

MODEL MANAGER

Purpose and functionality: A model deployment tool that integrates models into a business process addresses and solves a lack of visibility into the model development process. Without visibility, it is impossible to determine the stage of development and who is working with the model at each point as it goes through the analytic lifecycle. Lack of visibility can lead to inconsistencies, greater scrutiny, and an inordinate amount of time and resources required to put models into production. Models must be validated in both test and production environments. They must be monitored for change and effectiveness. Reports must be created for planning, governance, and auditing. Time must be allotted for re-training if model performance drops below a configured threshold since over time, model performance often degrades. Workflow inside of Model Manager provides consistency and manages alerts. For example, a statistician might receive an alert to investigate the model results in production or a data scientist might receive an alert to determine whether anything in the data has changed. A modernization effort can implement consistent and automated model validation and deployment processes and can shorten the time required to put models into production. Model reports help maintain compliance with regulatory controls by providing information and documentation. The reports can preserve IP that might be lost due to the retirement or job transfers of experts. Models can be published to SAS®, DB2, Greenplum, Hadoop, Netezza, Oracle, SAP HANA, and Teradata.

Target persona: Advanced business analysts business process stewards.

Display 2 is an overview of the process flow using Model Manager as part of the overall analytics lifecycle.



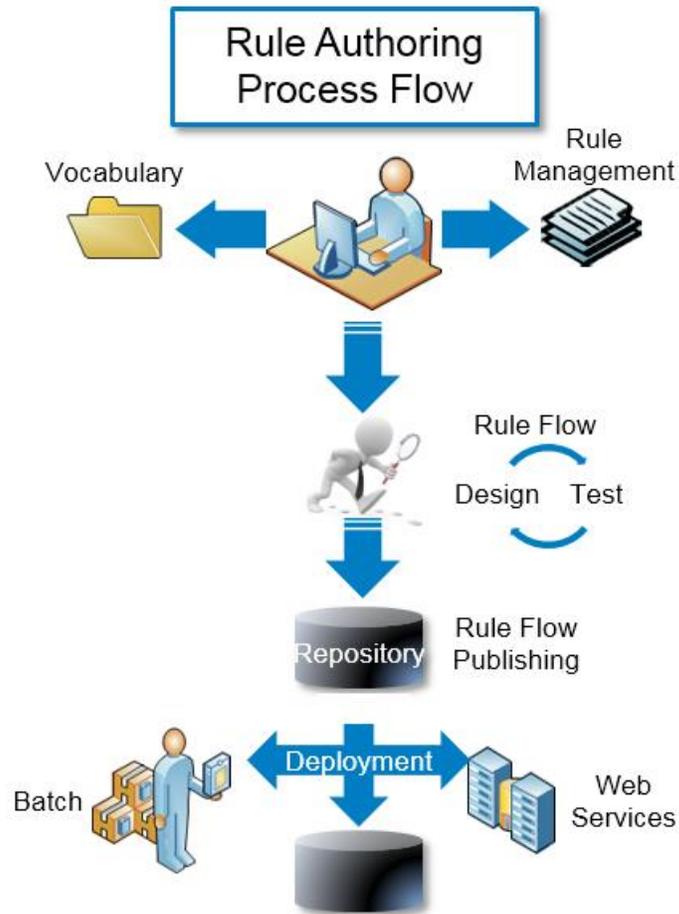
Display 2. Model Manager Process Flow

BUSINESS RULES MANAGER

Purpose and functionality: SAS Business Rules Manager provides a non-programming development environment for the following tasks:

- creating and managing vocabularies
- authoring and discovering rules and reusable rule sets that have been developed
- testing, updating, and publishing rules for production
- using an integrated workflow to version and govern rule flows comprised of rule sets

Display 3 shows an overview of a rule authoring process flow using Business Rules Manager.



Display 3. Business Rule Manager Authoring Process Flow

This section contains an explanation of business rules in this context, types of rules, rules for decisions, the process of managing rules as corporate assets, capabilities of the Business Rules Manager, and personas involved in an overall business process.

A business rule is a statement of business logic that specifies conditions to be evaluated and actions to be taken if those conditions are satisfied. “Business rules can vary a lot by industry, but understanding the need for business rules in decision processes is common across every application. ... Time is spent focusing on vocabulary creation – something that is necessary before any business rules are written.

Vocabularies need to be recognizable to the business and must be organized in a way that reflects the natural logic of the business application. A data driven approach to selectively creating or deriving a vocabulary has been very useful and well received by customers in every industry.

... In its most common form, a vocabulary is a ‘fundamental tool for communication and acquiring knowledge,’ according to Wikipedia. In the context of decision management, a vocabulary is a set of terms that are used to create business rules. It is part of the business rules language, along with operators that define how the terms are combined.”²

FORMS OR TYPES OF RULES

The phrase “business rules” is generally used loosely and can refer to constraints in a query, a data mapping, a data quality constraint, a data transformation, a model, a policy enforced for a business process, and so on.

“Business rules are simple expressions that analysts typically enter into a modeling process to support a defined business strategy. For example, you might decide that in addition to considering the customer’s probability score, you want customers in a given region to get a certain type of loan, or customers of a certain demographic to get a specific promotional offer”.⁴

Rules enable decisions. First, “we separate decisions from processes, systems, and events so they can be identified and managed explicitly.” Then “we manage the rules for those decisions as coherent rule sets, each with a singular focus and common purpose. The management of these rule sets is handled by those who are responsible for the source of those rules – legal documents, company policy, marketing strategy, call center expertise, and so on.”⁵

Rules are often managed as corporate assets. Business rules platforms are a key enabler of today's architectural imperative to build for constant change. The platform gives business analysts the authoring tools to maintain decision logic that is usually embedded deep in application code. It also provides a foundational vocabulary and collaborative IT processes to support the work of analysts.

Purpose and functionality: Support business persona with a business rule development interface that does not require programming skills. Business users and analysts who fully understand the business can author, test, manage, and apply decisions and business rules to determine next actions. They then work with IT to deploy the business rules in production.

Within Business Rules Manager, the business persona can create and manage vocabularies and can author, manage, and publish rules. A rule set is a logical collection of rules. A single rule set can have many rules. A rule flow is a logical collection of rule sets. A rule flow defines a set of rule sets and the order in which they must be executed. Frequently, a single rule flow corresponds to a single decision.

BUSINESS PERSONA EXAMPLES

Rules Designer: Domain Expert, Non-technical

Designs and reviews all rules that will be put into production. Validates the accuracy of rule types and behaviors, and categorization or grouping. Defines all business terms that will be used for authoring business rules. This person researches and validates rule names and logic from legal documents, company policy, marketing strategy, call center expertise, current business processes, excel spreadsheets, manual processes, and so on.

Rules Vocabulary Owner and Rules Author: Domain Expert, Semi-technical Business User

Works with the Rules Designer and with documentation to create vocabulary, entities, and terms. Creates lookup tables. Authors rules per design specifications, tests rules, modifies rules, and retests.

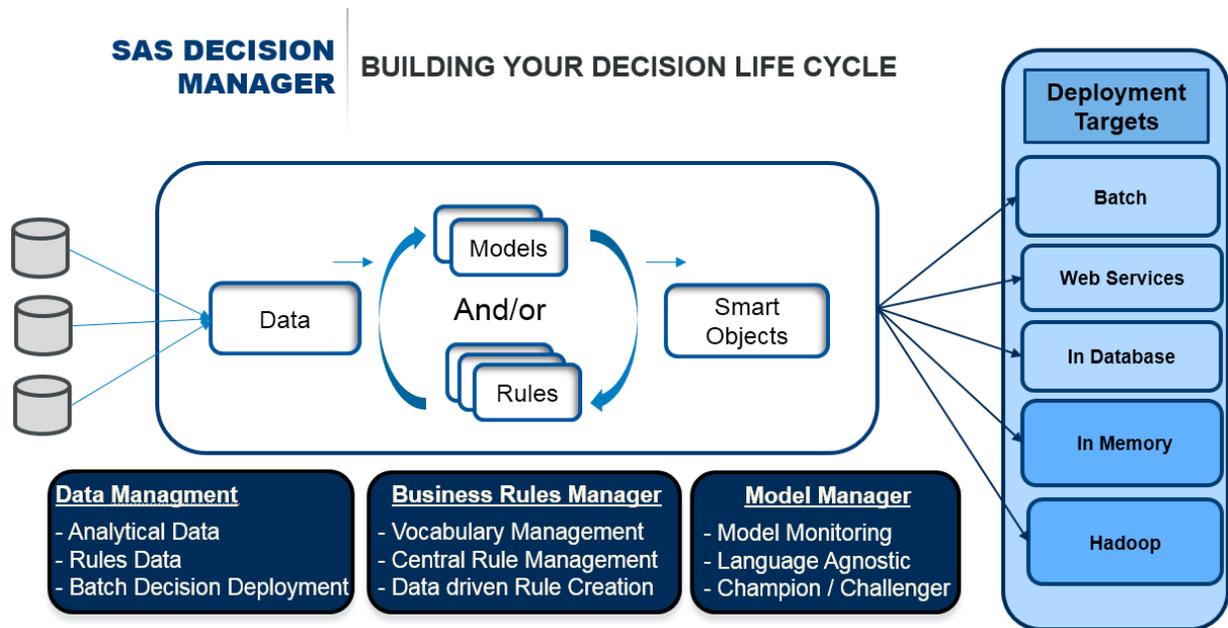
Rules Steward: Domain Expert, Business Process Experience, Semi-technical

Oversees the design and development of the rules. Is responsible for publishing and documenting rules and versions and keeping track of the applications on which rules are deployed. Works with the Rules Author.

DATA MANAGEMENT FOR RULES, MODELS AND DEPLOYMENT

OVERVIEW

Display 4 shows an overview of a generalized deployment process flow.



Display 4. Generalized Decision Life Cycle

INTEGRATION WITH ANALYTICAL MODELS

Linking business rules and analytical models through shared metadata and operational data at deployment allows an organization to control, manage, and automate the repeatable decisions at the heart of its business.⁴

DEPLOYMENT PERSONA EXAMPLE

Deployment Steward: Technical, Business Power User or IT

Works with the Rules Steward and connects data and rules/models within the batch user interface, SAS® Data Integration Studio. Tests results and serves as a liaison with IT to deploy rules/models to production environments. Deployment options are the following: batch, inDatabase to Teradata or Greenplum; SOAP or RESTful web services.

PUTTING IT ALL TOGETHER WITH SAS DECISION MANAGER

USE CASE: AUTO AUCTION CAR SALES FROM A BUYER'S PERSPECTIVE

Each year, hundreds of thousands of used cars are sold at auto auctions. The auction car buyer needs to make the best decision for each car and for the entire portfolio of purchases. Several factors influence the decision on purchasing cars at auction:

- Auto Market: what vehicles are popular, for what buyers, and in what regions
- Pipeline: which vehicles are in surplus or deficit on the buyer's car lots
- Strategy: what is the buyer's resale strategy
- Vehicle Condition: which cars boost the buyer's reputation when reselling to his or her customers
- Vehicle Cost: what is the expected purchase cost and the resale value

Assume that the following tasks have been completed.

- Data to support the modeling effort and business rules has been registered in metadata.
- The model has been designed and built using Enterprise Miner. (SAS STAT could also have been used for development and testing.) The model package has been exported from Enterprise Miner.

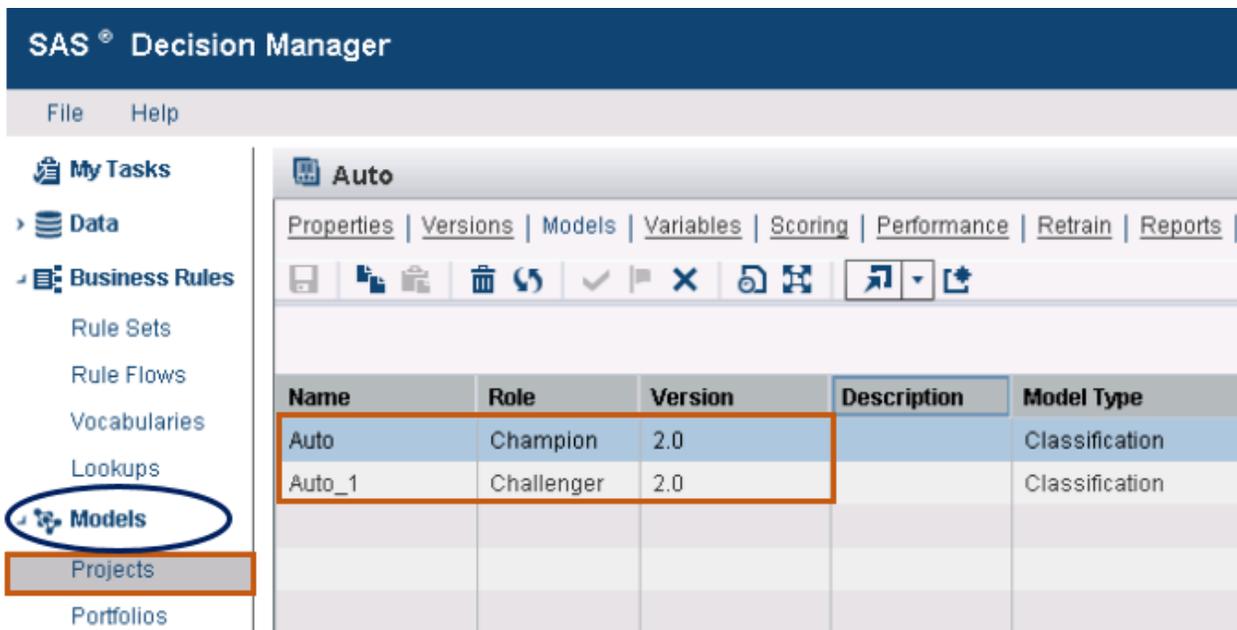
- Model Manager has been used to import the model package and we can see the properties of all the models. The champion and challenger models have been selected and published to SAS metadata via Model Manager.
- “The Auto model can capture hundreds or thousands of variables from the data, but you will probably also want to add more logic that reflects your unique business policies, customer base, or market conditions.”⁴ The auto model might be adjusted to include additional information that is unique to a particular area of the country or lifestyle.

Tasks shown in this paper are using modules in the SAS Decision Manager environment that enable customers to build, test, and publish business rules. The tasks are listed here:

- Build out one rule flow and test using SAS Business Rules Manager. This rule flow is used in conjunction with the Auto champion model. There are four rule sets in the one rule flow. We show one rule set as an example.
- Test the business rules against the model output or scored data.
- Publish the rule flow to metadata.
- Combine data, business rules, and predictive models into one auto auction purchase decision strategy. This will be a batch deployment of the champion Auto model linked with business rules in SAS Data Integration Studio.
- Deploy web services via SAS BI Web Services and test in SoapUI. Use a third-party web interface to show how the auction buyer can access the decision process in near real time from the auction site.

In SAS Decision Manager, we see the Auto Project Folder as shown in Display 5. The view shown is Model Manager in the SAS Decision Manager container. The Auto model has been set as the Champion and is the published model for consumption by other applications. There are two models in the Auto project and the model named Auto_1 is the challenger. There can be one or many challengers. In the Model Manager Process Flow shown in Display 2, we are at the deploy stage and we are ready to build and test business rules that will be linked with the model output at batch deployment in SAS Data Integration Studio.

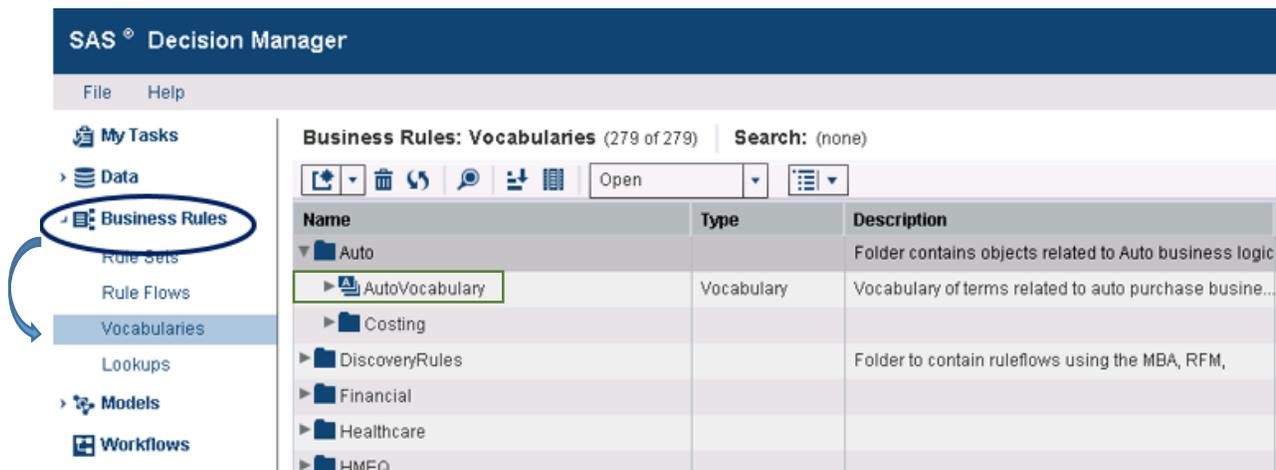
Display 5 shows Model Manager with one project named Auto. The Auto project contains two models, Auto and Auto_1.



Display 5. Model Manager with Two Models

BUSINESS RULES MANAGER

In SAS Decision Manager, we see the AutoVocabulary vocabulary in Display 6. The view shown is SAS Business Rules Manager in the SAS Decision Manager container.



Display 6. Business Rules Manager

Within the Auto vocabulary in Business Rules Manager, we have several entities, as shown in Display 7. You can use entities to organize or group similar business application domains.

Name	Type
▼ Auto	
▼ AutoVocabulary	Vocabulary
▶ AutoActions	Entity
▶ AutoTopEntity	Entity
▶ CarScores	Entity
▶ CLASS	Entity
▶ Costing	Entity

Display 7. Entities in the AutoVocabulary Vocabulary

Display 8 shows the vocabulary terms. These terms are mapped to actual data column names during testing in Business Rules Manager and during deployment in Data Integration Studio.

Name	Type	Description
▼ Auto		Folder contains objects related to Auto business logi
▼ AutoVocabulary	Vocabulary	Vocabulary of terms related to auto purchase busine
▼ AutoActions	Entity	Entity for all business rule action terms
123 AuctionPriceRatio	Term	price ratio checks
123 AverageCarCost_BYGROUP	Term	
BadCar	Term	final check indicator if vehicle is good or bad
BadMakeIndicator	Term	Flag for pass or fail of MAKE rules
▲ BadMakeReason	Term	
BadStateIndicator	Term	indicates if VIN State is valid or invalid, T or F
▲ Collector	Term	
123 Count	Term	Counter for cars
▲ Flag	Term	
▲ InvalidCarMake	Term	
▲ Lookup_Car_Make_Value	Term	
▲ Lookup_Check	Term	
▲ Lookup_Verify	Term	
▲ make	Term	
123 MileagePerYear	Term	average miles per year
MileageRatioBad	Term	
PriceRatioBad	Term	This price ratio is too high for business

Display 8. Vocabulary, Entity, and Terms

As shown in Display 9, a rule set called FinalBadCarCheck has been created. The action term, BadCar has been set to True or False. This value depicts the overall decision of whether buying a particular car at auction is a good buy or not. Notice we are using the CarScore condition term, which relates back to the

EM_PROBABILITY output term from the Auto model managed in Model Manager. In this way we are adding logic to what the model outcome would be and potentially overriding model output with better information found in the business rules.

Display 9 shows conditions and actions within the Business Rules Manager.

Vocabularies		Condition Term				Action Term
#		CarScore	MileageRati...	BadStateIn...	wheeltypes	BadCar
1	If				Initialized to False	False
2	If	< .5	True			True
3	If	< .5		True		True
4	If	< .5			'Special'	True
5	If	> .5				True

Display 9. Conditions and Actions within the Business Rules Manager

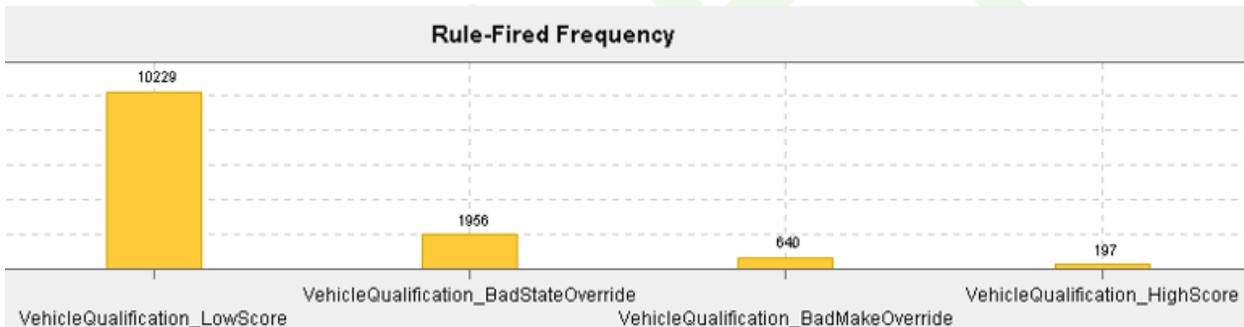
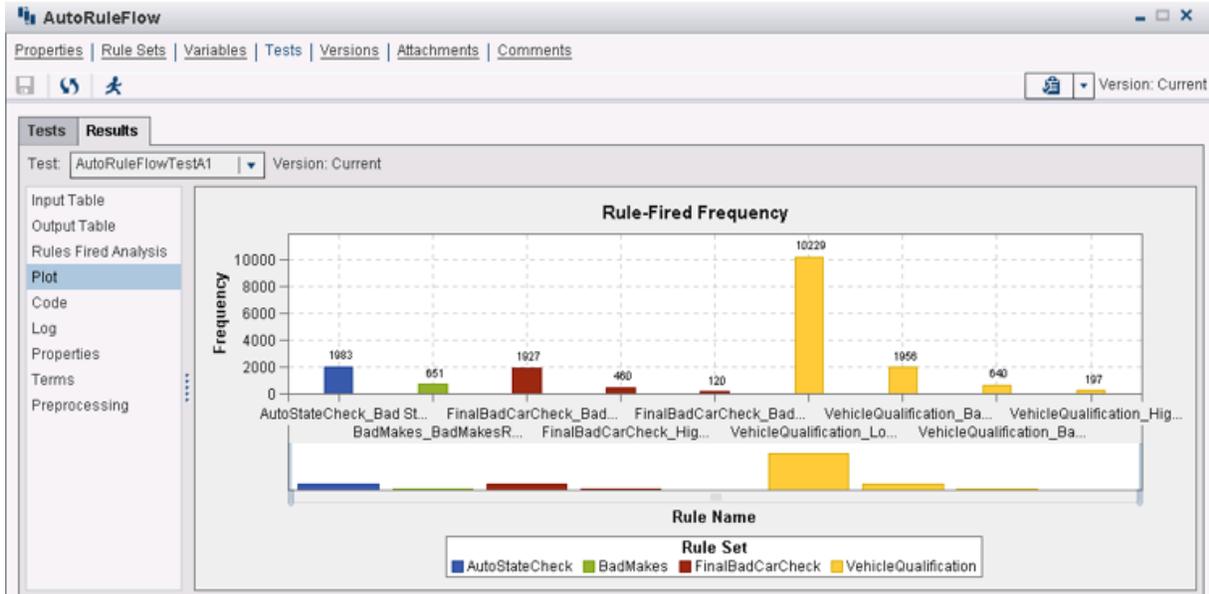
As shown in Display 10, a rule flow called AutoRuleFlow with four rule sets has been created. Each rule set contains one or more rules or rows. Rule Set 4, FinalBadCarCheck, is highlighted.

AutoRuleFlow					
Properties Rule Sets Variables Tests Versions Attachments Comments					
Simple Rule Flow					
Rule Sets					
#	Run	Name	Version	Description	
1	<input checked="" type="checkbox"/>	AutoStateCheck	1.2	Rules for determining if vehicle's VIN state is valid	
2	<input checked="" type="checkbox"/>	BadMakes	1.4	Identify Vehicles with bad makes	
3	<input checked="" type="checkbox"/>	VehicleQualification	1.2	Rules for Vehicle Qualification	
4	<input checked="" type="checkbox"/>	FinalBadCarCheck	1.7	uses other checks and car scores from model to d	

FinalBadCarCheck: Version 1.7 Rules Logic (5 Rules)			
#	Rule Name	Condition	Action
1	InitBadCar	If	then BadCar = False
2	MileageRatioBad	If (CarScore < .5) AND (MileageRatioBad = True)	then BadCar = True
3	BadState	If (CarScore < .5) AND (BadStateIndicator = True)	then BadCar = True
4	BadWheelType	If (CarScore < .5) AND (wheeltypes = 'Special')	then BadCar = True
5	HighScore	If (CarScore > .5)	then BadCar = True

Display 10. Rule Flow with Four Rule Sets

Test the rule flow in Business Rules Manager by pointing to a data table that has been registered in metadata for Business Rules Manager to access. It is easy to see how our rules fired using a plot within the Business Rules Manager interface. As shown in Display 11, the FinalBadCarCheck rule set has four rules that have fired 10229, 1956, 640, and 197 times respectively. The legend colors relate to the rule set name. Each rule set contains one or more rules.



Display 11. Rule Fire Graph

After testing and observing the rule fire behavior and examining our results, we can publish the rule flow to the business rules repository and it will be available via metadata to other applications.

Display 12 shows how to publish in metadata for consumption by other applications.

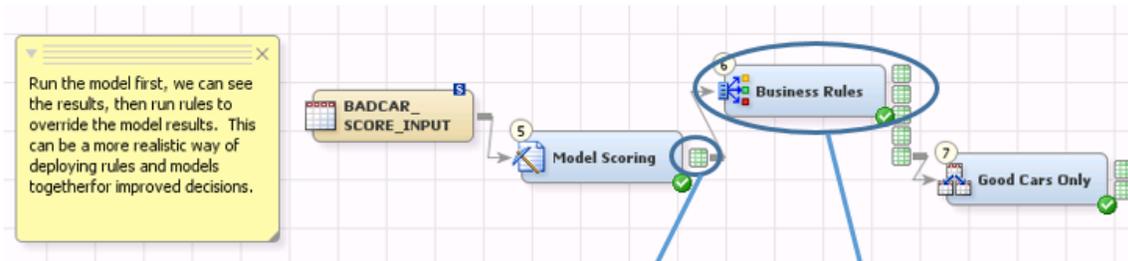
#	Publish rule flow	Version	Description
1	<input checked="" type="checkbox"/>	AutoStateCheck 1.2	Rules for determinin
2	<input checked="" type="checkbox"/>	BadMakes 1.4	Identify Vehicles with
3	<input checked="" type="checkbox"/>	VehicleQualification 1.2	Rules for Vehicle Qu
4	<input checked="" type="checkbox"/>	FinalBadCarCheck 1.7	uses other checks a

Display 12. Rule Flow Publish Indicator

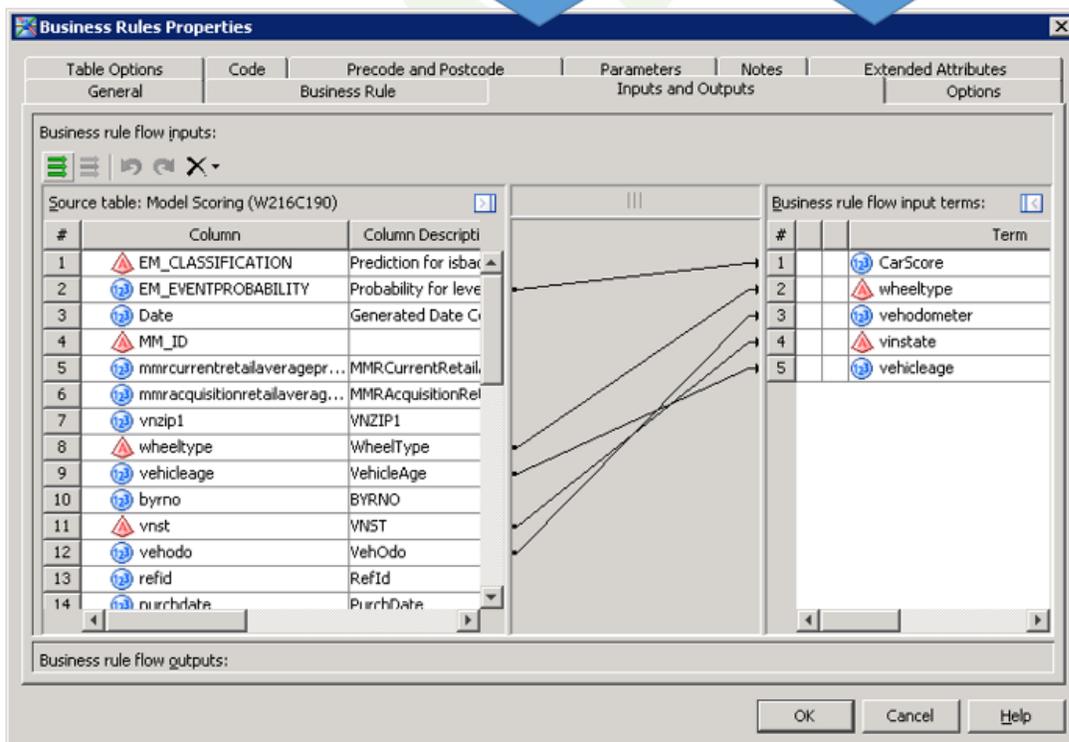
SAS Data Integration Studio is the interface that provides a batch decision flow builder, scheduling, and deployment environment where data, models, and rules can be linked and put into production. Data management needs that are outside of the normal customer IT processes and are approved by IT for usage in the overall business process for implementing models and rules to production environments are part of the data management processes prior to actual execution of the rules and models.

We can run the Auto model against the input table and then run business rules on the model output. This process links models and rules. The rules can override the model response based on external conditions such as the U.S. state that the car came from, the model of the car, age, transmission, and wheel type. The final outputs are decisions based on the model and the rules together. This approach yields an improved decision over running the model by itself.

Display 13 shows linking models and rules within the batch decision flow builder in Data Integration Studio.



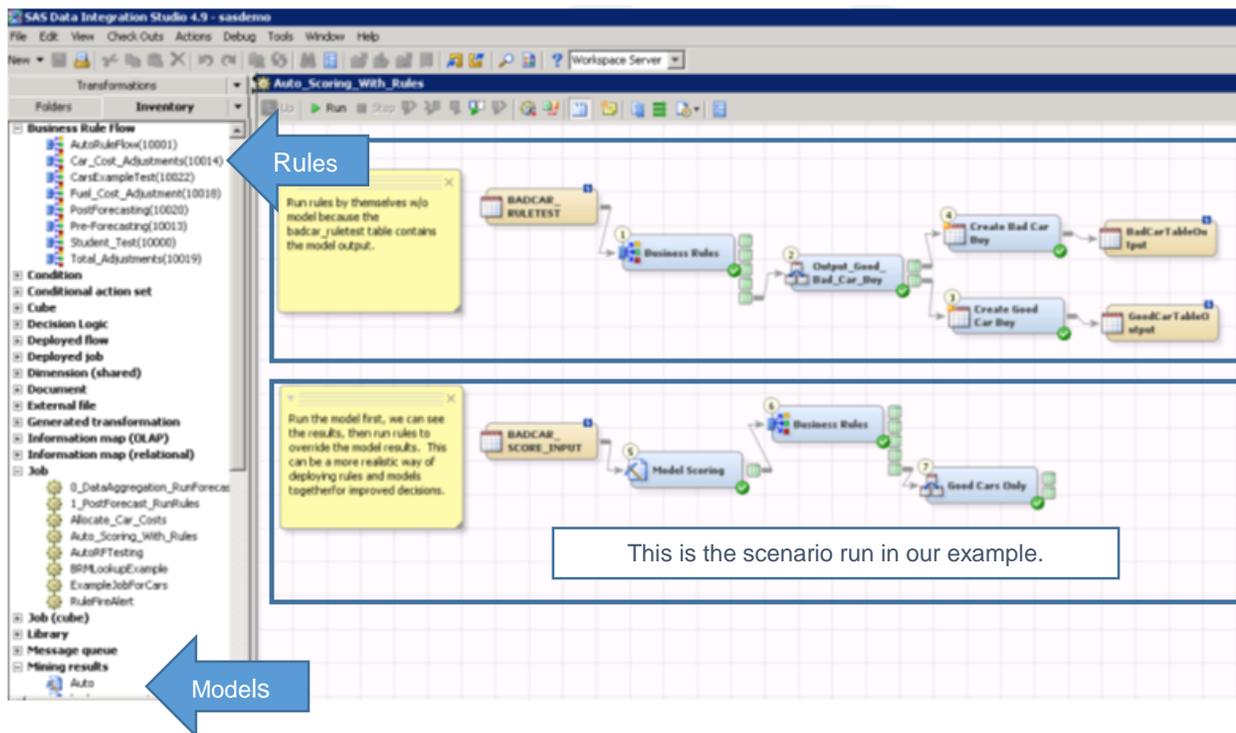
Business Rules node, the model output table columns are mapped to rule flow input terms; column names and other BRM objects must obey SAS name requirements



Display 13. Business Rule Flow Input Terms on Right Side Mapped to Model Output

Batch deployment using SAS Data Integration Studio is shown in Display 14 below. We can run rules against data without models as seen in the top rectangle, or we can link models and rules together as shown in the lower rectangle. On the left side, we see the smart objects for business rule flows and

models, which are available for consumption within Data Integration Studio. Display 14 shows a Data Integration Studio job flow with rules and models linked.



Display 14. Data Integration Studio Job Flow with Data, Models, and Rules

The final output table shown in Display 15 results from running the scenario in the lower rectangle shown in Display 14. These are the results after running business rules on the model output data. The car buyer can use these results to determine which cars are most likely good buys. Because the business rules were run on the model output data, the decision results highlighted in yellow take into consideration mileage, state, and other external factors. Applying business rules to the model output data can lead to an improved decision over running the model without business rules. This data suggests that the car buyer should be wary of buying certain cars from Texas.

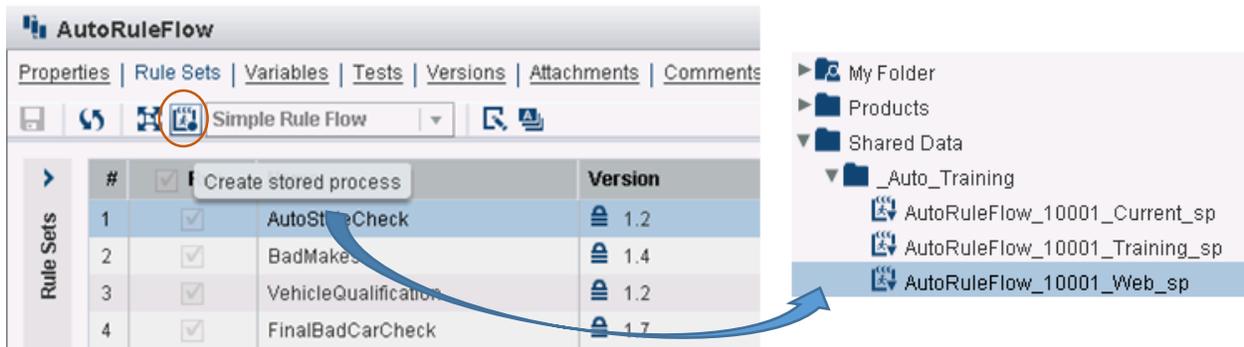
Note: Not all car attributes are shown in the table below in Display 15.

#	vinstate	CarScore	BadCar	EM_CLASSIFICATION	vehicleage	vehodometer	MileagePerYear
1	FL	0.0632370882	0	0	4	85589	21397.25
2	NC	0.0632370882	0	0	3	62279	20759.666667
3	AZ	0.1503919007	0	0	5	81739	16347.8
4	NC	0.1503919007	0	0	6	73560	12260
5	CA	0.0632370882	0	0	3	52381	17460.333333
6	OK	0.0632370882	0	0	4	88004	22001
7	NC	0.0632370882	0	0	2	83346	41673
8	TX	0.0632370882	1	0	3	77656	25885.333333
9	TX	0.0632370882	1	0	2	55866	27933
10	LA	0.1503919007	0	0	5	86450	17290
11	TX	0.0632370882	1	0	3	77652	25884

Display 15. Final Output after Running Rule Logic against Model Output

You can also deploy business rule flows using a SOAP web service after creating a SAS stored process. As shown in Display 16, one click from the Business Rules Manager creates a stored process. The user is then asked for a metadata folder location for organization.

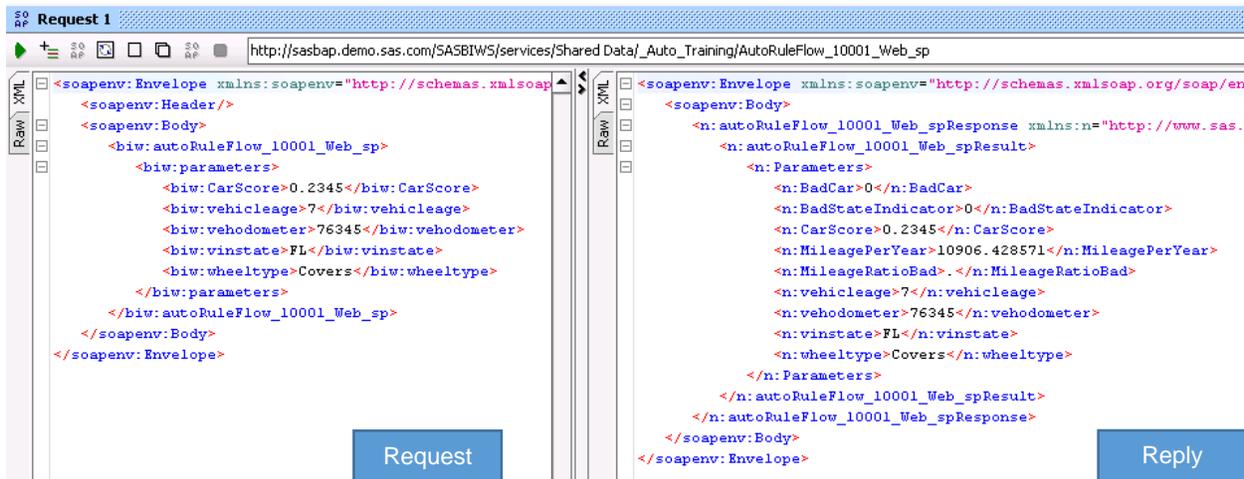
Display 16 shows how to create a SAS stored process for the AutoRuleFlow rule flow by clicking the icon indicated within the orange circle. The stored process is named AutoRuleFlow_10001_Web_sp and is saved in the metadata folder location chosen by the user.



Display 16. Create a Stored Process and Save It in a Metadata Folder

After the stored process is registered in metadata, it is immediately available as a SOAP web service and can be tested with source SoapUI or some other web service testing facility, as seen below in Displays 17 and 18.

Display 17 shows the request and response of the AutoRuleFlow_10001_Web_sp web service in SoapUI.



Display 17. SoapUI Shows Left Side Request and Right Side Reply

Display 18 shows an example of a third-party interface calling the AutoRuleFlow SOAP web service.

The screenshot shows the SAS Web Service Client interface. At the top, it says "SAS Web Service Client" with a home icon and a close button. Below that, there are tabs for "Request" and "Response". The "Request" tab is selected, and it shows a "REST" method. Below the request tab is a table with two columns: "Variable" and "Value". The request table contains the following data:

Variable	Value
Car Score	0.23707
Vehicle Age	8
Odometer Reading	74634
VIN State	TX
Wheel Type	Covers

Below the request table is a "Send Request" button. To the right of the request tab is the "Response" tab, which shows a response time of "8559 ms". Below the response tab is a table with two columns: "Variable" and "Value". The response table contains the following data:

Variable	Value
Bad Car Buy Decision	1
Bad State Indicator	1

Display 18. External Web Application That Calls the AutoRuleFlow Web Service

CONCLUSION

Modernizing business by more pervasively applying analytics is a compelling goal. However, it takes time and planning to implement a solution. With the ever increasing volumes and types of data and the growing speed or frequency at which this data is generated, the data, apart from the analytics, should be well managed. The data must support reliable and relevant models, which can be deployed at various decision points within an organization or business. Building advanced statistical and mathematical models off of this data provides necessary trust of the outcomes that will be used for decision making. Because highly skilled individuals are valuable resources, being able to optimize skills and provide opportunities for skill set growth, businesses require tools and solutions that enable self-service capabilities. These self-service capabilities are translated into collaboration with IT in terms of development and deployment. Thus, the overall business process has improved efficiency and decisions based on rapid and robust applications of analytics into operational systems. This is a repeatable established process, which can provide the necessary edge on competition, customer-specific interactions, responding to ever changing business requirements and governmental regulations.

REFERENCES

- ¹Bloomberg Businessweek Research Services, 2011. SAS Institute white paper. "The Current State of Business Analytics: Where Do We Go From Here?" Available <http://www.sas.com/reg/gen/corp/1504633>.
- ²Crain, Charlotte. SAS blog. "Vocabularies: Deciphering the language of business rules." Available <http://blogs.sas.com/content/sascom/2014/12/17/vocabularies-deciphering-the-language-of-business-rules/>.
- ³Duling, David, Steve Sparano, and Charlotte Crain. 2014. "Introducing SAS Decision Management." *Proceedings of the SAS Global 2013 Conference*. Cary, NC. SAS Institute Inc. Available <http://support.sas.com/resources/papers/proceedings14/SAS276-2014.pdf>.
- ⁴Duling, David. ; SAS Conclusions Paper. "How to Implement an Effective Decision Management System." Available http://www.sas.com/en_us/whitepapers/how-implement-effective-decision-management-system-106220.html.
- ⁵Taylor, James. JT on EDM blog, March 5, 2009. "Here's how decisions and rules relate and how to manage them." Available <http://jtonedm.com/2009/03/05/heres-how-decisions-and-rules-relate-and-how-to-manage-them/>.

RECOMMENDED READING

- "How to Manage Your Data as a Strategic Information Asset." Available <http://sas.com/reg/wp/corp/47976>.

- “Time Is Precious, So Are Your Models: SAS® Provides Solutions to Streamline Deployment.” Available <https://support.sas.com/resources/papers/proceedings13/086-2013.pdf>.

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the authors at:

Charlotte Crain

Charlotte.crain@sas.com

Steve Sparano

Steve.sparano@sas.com

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are trademarks of their respective companies.