

4

Principles of Decision Management Systems

Organizations that have adopted Decision Management Systems have gained tremendous results from doing so. The use of business in Decision Management Systems has given organizations the agility to respond rapidly to competitive and market changes, to avoid business risks, and to take advantage of narrow windows of opportunity. The use of analytics to predict risk, fraud, and opportunity in these Decision Management Systems has kept companies profitable despite the risks they face and has allowed them to maximize the value of their customer relationships through a laser focus on opportunity. The ability of Decision Management systems to adapt to change and to be part of a learning environment has allowed organizations to experiment with new approaches, learn from their successes and failures, and continuously improve their business. Any organization would want the benefits of these kinds of systems.

However, it is not immediately clear how to build Decision Management Systems. Although there are specific technologies involved, the use of these technologies is not sufficient to ensure that Decision Management Systems are the outcome of using them. Decision Management Systems appear to deal with different issues, and have different characteristics, across different industries and business functions. It can be hard to see what an underwriting system in use by the agents of a property and casualty insurer has in common with a real-time offer management system supporting a website. Yet when the basic principles of Decision Management Systems are

understood, they can be correctly identified and delivered with maximum return on investment.

Four specific principles are at the heart of identifying and building Decision Management Systems. If a system exists, it can be assessed against these four principles and can be said to be a Decision Management System if these principles guide its design and implementation. If an information systems project is being considered, then the integration of these principles into the project will ensure that what is delivered is a Decision Management System.

The four principles address the characteristic capabilities of a Decision Management System:

1. Begin with the decision in mind.

Decision Management Systems are built around a central and ongoing focus on automating decisions, particularly operational and "micro" decisions.

2. Be transparent and agile.

The way Decision Management Systems make each decision is both explicable to non-technical professionals and easy to change.

3. Be predictive, not reactive.

Decision Management Systems use the data an organization has collected or can access to improve the way decisions are being made by predicting the likely outcome of a decision and of doing nothing.

4. Test, learn, and continually improve.

The decision-making in Decision Management Systems is dynamic and change is to be expected. The way a decision is made must be continually challenged and re-assessed so that it can learn what works and adapt to work better.

Principle #1: Begin with the Decision in Mind

Most information systems have been developed, and are continuing to be developed, around business functions, business data, or business processes. Functional systems support a set of related business functions such as accounting or human resources. Data-centric systems focus on particular kinds of data, such as customer information. Business processes such as order-to-cash have been layered on top of both kinds of systems, and newer systems are developed to deliver additional specific business processes. Each of these

approaches has pros and cons, but they all share a common challenge—they assume either that people will make all the decisions involved in the functions and business processes being automated or that how these decisions are made can be fixed. As a result, none of them are Decision Management Systems.

To develop Decision Management Systems, we must take a different approach—not one based on functions, data, or processes. To develop Decision Management Systems we must begin with the decision in mind. Decision Management Systems are built to automate and improve specific business decisions. As a decision involves making a selection from a range of alternatives, these systems make that selection—they choose the action or actions that can or should be made given the data available and the context of the decision. Decision Management Systems do not assume that every decision must always be taken by a human. Decision Management Systems make these decisions using the same business logic humans would apply without human intervention.

Clearly, however, we are neither willing nor able to build information systems to make every decision on our behalf. Only certain decisions can and should be addressed by Decision Management Systems.

What Kinds of Decisions Are We Talking About?

Information systems are good at handling repetitive tasks. They excel at doing the same thing over and over without variation and without making mistakes from one transaction to the next. Something that cannot be defined in a repeatable way is not a good target for any kind of information system; thus, only those decisions that are repeatable are good candidates for being automated and managed using a Decision Management System.

Repeatable Decisions

A repeatable decision is one that is made more than once by an organization following a well-defined, or at least definable, decision-making approach. Business decisions can be categorized in various ways; one effective way to look at decisions is to categorize them as strategic, tactical, or operational (Taylor & Raden, 2007). This divides decisions into three categories based on the value of each decision made—the difference between a good and a bad decision—and the number of times such a decision is made by an organization:

- √ Strategic decisions are those high-value, low-volume decisions that guide the overall direction of the company. These ad-hoc, typically one-off decisions are made by senior management or the executive team of an organization. Lots of information is assembled and analyzed while many options are considered. Once the decision is made, it is never made again in the same context—even if it is revisited later, this is really a different decision as circumstances are different. Organizations may know that a strategic decision is going to be needed well in advance, but often these decisions arise from unexpected opportunities or challenges. Strategic decisions are not candidates for Decision Management Systems as they lack the key element of repeatability.
- √ Tactical decisions are those focused on management and control. These medium-value decisions still have significant business impact. They too involve data and analysis, typically by humans in management or knowledge worker positions. However, these decisions do repeat—the same kind of decision is made repeatedly during normal business operations. Decisions about the discounting approach being used or the staffing levels of a call center might be examples, and these decisions must be made every month or every week. The same or very similar analysis is performed each time, and company policies may play a significant role in how the decision is made. More repeatable and consistent tactical decisions are certainly targets for Decision Management Systems.
- √ Operational decisions are those of lower individual value and typically relate to a single customer or a single transaction. They are critical to the effective operation of an organization, especially an organization of any size. Because of the number of times they must be made, consistency and repeatability are critical. Policies and well-defined decision making criteria are typically developed to ensure this consistency. Despite their low individual value, they are extremely valuable in aggregate. A decision made thousands or millions or even billions of times a year has a total value that often exceeds even the most important strategic decision.

Furthermore, strategic and tactical decisions (for example, to focus on customer retention or discount more aggressively) will only have an impact if a whole series of operational decisions (how to retain this customer or what discount to offer this distributor) are made in accordance with the higher-level decision. For these reasons, operational decisions are the most common subject of Decision Management Systems.

To begin with the decision in mind, we must understand what operational or tactical decision is to be the focus of our Decision Management System.

BUSINESS STRATEGY AND STRATEGIC DECISIONS

This focus on repeatable operational and tactical decisions can and should be combined with a focus on business strategy. A business strategy must be supported by many operational and tactical decisions if it is to be put into practice. For instance, if a focus on growing per-customer revenue is central to your business strategy, then there may be strategic decisions that must be made to support this strategy. There will definitely be many operational and tactical decisions that will be influenced by and contribute to this strategy. For example, unless operational decisions about customer retention and cross-sell offers are made effectively, you cannot deliver on this customer-centric strategy. As discussed in Chapter 5, "Discover and Model Decisions," the right operational decisions to focus on are those that support the objectives and key metrics of the organization.

Operational Decisions

Operational decisions are by far the most common kind of repeatable decision. Every order placed, every customer interaction, every claim, or credit card transaction involves operational decisions. Operational decisions are the day-to-day, run-the-business decisions that are taken in large numbers by every organization.

Operational decisions are highly repeatable—in fact, being consistent by following a set of guidelines or applying the relevant policies and regulations is a defining characteristic of an operational decision. Operational decisions can also involve an assessment of risk, as many forms of risk (loan default or credit

risk, for instance) are acquired one transaction at a time. Operational decisions often must be made in real time or near real time, while customers are waiting for the decision to be made.

Although many operational decisions are made about customers, they can also be made about shipments, suppliers, or staff. As more physical devices are connected to the Internet with sensors or RFID chips, operational decisions are often made about "things"—about vehicles, packages, railcars, or network components.

Micro Decisions

Micro decisions are a particular kind of operational decision (Taylor & Raden, 2007) where the desire to personalize an interaction with a customer requires a focus on making a decision for that customer and that customer only. Often an operational decision is repeated for all customers, with the decision being based only on the data available for the particular interaction or transaction concerned. A micro decision, in contrast, uses everything known or predictable about a customer to make a unique decision just for them. Two customers making the same request or involved in identical transactions would get two different outcomes.

This focus on the information about the customer is what makes micro decisions a distinct form of operational decision. Everything known about the customer must be synthesized into actionable insight about the customer and fed into the operational decision alongside the information about the transaction. For example, when an order is placed, two operational decisions might be made—what shipping options to offer on the order and what discount to offer. The first of these might be managed as a standard operational decision with information about the order such as delivery address, weight, and value used to determine which of the various shipping options would be allowed. The second could be managed similarly but could also be handled as a micro decision. The customer's history with the company could be used to compute their likely future profitability and the risk that they might consider a competitor. This information, as well as information about the specific order, would then feed a micro decision to calculate a discount specific to this customer placing this order at this moment.

Tactical Decisions

Although most repeatable decisions are operational or micro decisions, some are more tactical in nature. These repeatable tactical decisions often relate to management control of operations, such as assessing the staffing level required by a call center for the coming shift. They may also include knowledge worker decisions, such as those in clinical situations where a doctor is advised as to the likely interactions of a set of medications she has just prescribed to a patient. Although these tactical decisions are not as high-volume as operational or micro decisions, they are often of slightly greater value and so still offer an opportunity for Decision Management Systems.

CHANGING DECISION CRITERIA

Another set of tactical decisions under managerial control are those for revising the decision criteria used in operational decisions. The world is dynamic, so the decision criteria for operational decisions need to change regularly. For example, customer preferences and fraud patterns change over time, as do the criteria for deciding what offers are to be made. Some tactical decisions are about setting the right decision criteria to be applied in an operational decision.

Because these systems are more complex and less repeatable, however, it is likely that the system will not completely automate the decision. Instead it will guide and support the decision maker by restricting the available options or by focusing them on a specific set of information, which will be useful to making the final determination.

Different Types of Decisions Interact

Operational decisions are made every time a business process or transaction executes, and tactical decisions are made periodically to change operational decision criteria or to attend to exceptional business situations and take corrective actions. Over a longer horizon, this is not sufficient to improve business outcomes. Business strategy guides tactical and operational decisions and may need to change to respond to the dynamic marketplace and the external world. For example, competitors may introduce new products and services, influencing customer choice and putting

competitive pressure on revenue and profitability. This may mean changing business strategy around providing targeted discounts and customizing products. This in turn creates change in processes and associated operational decisions. These strategic, tactical, and operational decisions must be aligned.

One way to understand the relationship between operational, tactical, and strategic decision-making is shown in Figure 4-1. This Observe-Orient-Decide-Act (OODA) model was originally introduced by military strategist and USAF Colonel John Boyd. Business outcomes are "observed" to detect changing situations that may lead to new tactical decisions, represented by "decide" or strategy "reorientation" with changes in business processes and additional decisions. "Act" represents operational decisions following the decision criteria set by the "decide" stage. There's more on the OODA Loop in Chapter 10, "Technology Enablers."

Figure 4-1 The Decision Lifecycle From strategy definition to decision automation.

If We Are Talking About Decisions, Aren't We Just Talking About Decision Support Systems?

The line between Decision Management Systems and Decision Support Systems (or Executive Information Systems) can be blurry. This is especially true when considering the kind of Decision Management System that handles tactical decisions, or where an operational decision is not completely automated—where the user is presented with multiple valid options, such as possible offers to make.

Decision Management Systems are distinct, however, and they differ from traditional Decision Support Systems in five ways:

1. Decision Support Systems provide information that describes the situation and perhaps historical trends so that humans can decide what to do and which actions to take. Decision Management Systems automate or recommend the actions that should be taken based on the information that is available at the time the decision is being made.
2. The policies, regulations, and best practices that determine the best action are embedded, at least in part, in a Decision Management System where a Decision Support System requires the user to remember them or look them up separately.

3. The information and insight presented in a Decision Support System is typically backward looking, and Decision Support Systems are generally reactive—helping human decision-makers react to a new or changed situation by presenting information that might help them make a decision. In contrast, Decision Management Systems use information to make predictions and aim to be proactive.
4. Learning is something that happens outside a Decision Support System and inside a Decision Management System. Users of Decision Support Systems are expected to learn what works and what does not work and to apply what they learn to future decisions. Decision Management Systems have experimentation or test-and-learn infrastructure built in so that the system itself learns what works and what does not.
5. Decision Management Systems are integrated into an organization's runtime environment. They make decisions for applications and services in the organization's enterprise application architecture. In contrast, Decision Support Systems are often desktop or interactive applications that execute outside the core application portfolio.

Why Don't the Other Approaches Work?

Before considering the remaining principles, it is worth considering why it is essential to begin with the decision in mind. What is it about a focus on functions, on process, or on data that prevents the effective development of Decision Management Systems?

A Functional Focus Is Not Enough

One traditional approach to building systems is to focus on a cluster of related functions—those to do with human resources or those to do with managing a factory, for instance. Such systems contain stacks of capability focused in one functional area and owned by a single functional department. This approach could result in the development of Decision Management Systems if the decisions involved were wholly contained within a single business function. However, while some decisions are concentrated in this way, many cut across functions. A discount calculation decision, for instance, might involve inputs from supply chain functions, from finance, and from customer management. As such a focus on functions will rarely identify and

encompass decisions in a way that lends itself to the construction of Decision Management Systems.

A Process Focus Is Not Enough

Functional applications have gradually fallen from favor as organizations have moved to focus on end-to-end business processes. Business processes such as "order to cash" or "issue policy" often cut across several functional areas, linking elements of one function together with elements of another to create a useful business outcome. Although this cross-functional approach can help with the identification of decisions, a pure process focus tends to entwine decisions with the process itself. If no real distinction is drawn between decisions and the processes that need those decisions, it is hard to create true Decision Management Systems. A strong separation of concerns, keeping business processes and decisions linked but separate, is required if enough of a focus on decisions is to be maintained.

Some processes keep decisions separate and manage them separately by assigning these decisions to people in manual process tasks. A focus on human decision-making, even in high-volume operational processes, also does not result in the construction of Decision Management Systems.

A Data Focus Is Not Enough

Particularly when constructing their own custom systems, organizations often focus on the data that must be managed. These systems become focused almost entirely on the management of the data elements or entities concerned. Providing what is known as CRUD functionality (Create, Read, Update, and Delete) for the core objects becomes their rationale. The data contained is managed only so that it can be edited and displayed while analysis is limited to reporting. Such systems often provide data for decision support systems but they, like process- and function-centric systems, defer decision making to actors outside the system.

Principle #2: Be Transparent and Agile

Most information systems in use today are opaque and hard to change. The use of programming languages—code to specify their behavior makes them opaque to any but the most technically adept. This opacity, and the difficulties of confirming that changes made to the code do what they are expected to do, make for long

change cycles and a lack of responsiveness. The combination means that extensive information technology projects must be planned, budgeted, and executed to make changes to the behavior of a system.

These characteristics are unacceptable in a Decision Management System. Opacity is unacceptable because many decisions must demonstrate that they are compliant with policies or regulations. If the code is opaque, then it will not be possible to see how decisions have been made and it will not be possible to verify that these decisions were compliant. Decision Management Systems also make decisions that are based on detailed business know-how and experience. If the code is so opaque that it cannot be understood by those who have this know-how or experience, then it is unlikely to be correct.

Organizational decision-making changes constantly, so agility is also essential. As regulations change, the behavior of any Decision Management System that implements that regulation must also change. Organizations also want Decision Management Systems to make good decisions—effective ones. Effective decisions based on the expectations of customers must be competitive, yet the behavior of competitors and customer expectations change constantly. And moreover, customers and competitors are not obliged to tell organizations when their expectations or plans change. An ability to rapidly change Decision Management Systems to respond is essential.

Decision Management Systems must therefore be both transparent and agile:

- √ The design must be transparent so that it is clear that the system is executing the behavior expected of it.
- √ The execution must be transparent so that it is clear how each decision was made.
- √ They must be agile so that their behavior can be changed when necessary without delay and without unnecessary expense.

Design Transparency and Why It Matters

A Decision Management System must exhibit design transparency. It must be possible for non-technical experts—those who understand the regulations or policies involved or who have the necessary know-how and experience—to determine whether the system is going to behave as required. Those without IT expertise must be able to manage the way in which decisions are made so that it is clear to all participants involved. The

drivers or source of this behavior must be identifiable so that those reviewing the behavior of the system can clearly assess its effectiveness in meeting objectives.

Tracking the source of decision-making behavior also means that changes in those sources can be quickly mapped to the changes required in the system. Design transparency means it is possible to determine the way in which a proposed change will ripple through the Decision Management System. One regulatory change might affect many decisions, for example, and decisions may be dependent on the same data elements because they have information needs in common.

Organizations must be sure that their Decision Management Systems will make decisions accurately and effectively after a change is made. This requires that the ripples and impacts of any change can be determined before it is made. Design transparency is essential to being able to trace these impacts.

Execution Transparency and Why It Matters

When a decision is made by a person, that person can be asked to explain the decision. If a person rejects an application for a loan, for instance, he can be asked to appear in court, to write a letter explaining, or simply to answer the customer's questions. This is not possible when a system makes a decision. A Decision Management System must therefore provide an explanation of a decision that will satisfy customers or suppliers who are materially affected by it. When a decision is regulated, such as when deciding which consumers may have access to credit, a Decision Management System must provide an exact description of how each decision was made so that it can be reviewed for compliance. Decision Management Systems must deliver real execution transparency in these cases.

Not all decisions require execution transparency. When marketing or promotional decisions are being made, it may not be necessary to understand exactly why a particular offer was made to a particular site visitor. When a Decision Management System is being used to decide when to bring a human into the loop, for fraud investigation for instance, it may also not be necessary to understand why as the human acts as a second "pair of eyes."

Even when a Decision Management System does not require execution transparency, an understanding of how each decision was made can help improve the decision-making of the system. Building in execution transparency is therefore generally a good idea,

whether it is required or not. Any approach to developing Decision Management Systems must support execution transparency as well as design transparency.

Business Agility and Why It Matters

An increase in transparency is likely to result in an increase in business agility—if it is easier to see how something works, it will be easier to change how it works when this is needed. A faster response to a needed change improves overall business agility. Transparency is necessary for agility but not sufficient. Once a change is identified and its design impact assessed, it must be possible to make the change quickly and reliably. Decision Management Systems can require real-time changes to their behavior in extreme cases. Daily or weekly changes are very common. When sudden market changes occur, such as major bankruptcies or an outbreak of hostilities, the resulting need for changes to Decision Management Systems can be extreme. Money—and perhaps lives—will be lost every minute until the change is made.

Decision Management Systems must change constantly to reflect new regulations, new policies, and new conditions. This rate of change must be both possible and cost-effective. For most businesses and other organizations, it will not be acceptable if the needed agility in Decision Management Systems comes at too high a price. For a Decision Management System, change must be easy, it must be reliable, it must be fast, and it must be cost-effective.

AGILE DECISION-MAKING FOR TRULY AGILE PROCESSES

Many organizations invest a great deal in developing agile business processes. Decision Management Systems further increase this agility as business changes often involve updates to business decisions. These decisions are often the most dynamic part of a process, the part that changes most often.

For instance, a company's pricing rules are likely to change far more often than its order-to-cash process. If only the business process can be changed quickly, then the company will not be able to respond to the far more numerous pricing changes without changing its process, an unnecessary step. Developing Decision Management Systems allows an organization to control business processes and the critical decisions within them. This increases the agility built into a process and allows

for a stable process even when decision-making is constantly changing and evolving.

Explicitly identifying decisions and describing the logic behind them allows this logic to be managed and updated separately from the process itself, dramatically increasing the agility of an organization.

Principle #3: Be Predictive, Not Reactive

In recent years, organizations have spent heavily on technology for managing and using data. Beginning with Database Management Systems and moving through Information Management, Data Quality and Data Integration to Reporting, and Dashboards, these investments are now mostly classified as Business Intelligence and Performance Management. These investments have taken data that was once hidden in transactional systems and made it accessible and usable by people making decisions in the organization.

These investments have been focused on analyzing the past and presenting this analysis to human users. They have relied, reasonably enough, on their human users to make extrapolations about the future. Users of these systems are making decisions based on this data, using what has happened in the past to guide how they will act in the future. Many of these systems can also bring users' attention to changes in data quickly to prompt decision-making. The value of this investment in terms of improved human decision-making is clear.

These approaches will not work for Decision Management Systems. When a decision is being automated in a Decision Management System, there is no human to do the extrapolation. Passing only historical data into a Decision Management System would be like driving with only the rear view mirror—every decision being made would be based on out-of-date and backward-looking data. In fact it would be worse, as a human driver can make guesses as to what's in front of her based on what she sees in a rear view mirror. She will be reasonably accurate too, unless the road is changing direction quickly. Systems are not that smart—without people to make extrapolations from data, Decision Management Systems need to be given those extrapolations explicitly. Without some view of the future and the likely impacts of different decision alternatives, a Decision Management System will fail to spot opportunities or threats in time to do anything about

them.

Predicting likely future behavior is at the core of using predictions in Decision Management Systems. You need to predict individual customer behavior such as how likely they are to default on a loan or respond to a particular offer. You need to predict if their behavior will be negative or positive in response to each possible action you could take, predicting how much additional revenue a customer might generate for each possible action. You want to know how likely it is that a transaction represents risky or fraudulent behavior. Ultimately you want to be able to predict the best possible action to take based on everything you know by considering the likely future behavior of a whole group of customers.

Decision Management Systems require predictions. They must be given predictions in the context of which they can act instead of simply reacting to the data available at the time a decision is made. They need access to predictions that turn the inherent uncertainty about the future into a usable probability. They cannot be told, for instance, which claims are definitely fraudulent—this is uncertain. They can be given a model that predicts how likely it is that a specific claim is fraudulent.

There are three specific ways in which Decision Management Systems can be given predictions. They can be given models that predict risk or fraud, that predict opportunity, and that predict the impact of decisions. They can use these predictions to direct, guide, or push decision-making in the right direction.

Predict Risk or Fraud

Most repeatable decisions do not have a huge economic impact individually. Despite their limited scope, many do have a significant gap between good and bad decisions. The value of the decision varies significantly with how well they are made. This gap arises when there is a risk of a real loss if a decision is made poorly. For instance, a well-judged loan offer to someone who will pay it back as agreed might net a bank a few tens of dollars in profit. A poorly judged offer will result in the loss of the loan principal—perhaps thousands of dollars. This mismatch between upside and downside is characteristic of risk-based decisions. Similarly, a poorly made decision in detecting fraud can result in large sums being transferred to an imposter or large purchases being made using stolen credit cards.

When Decision Management Systems are being used to manage these decisions, it is essential that the decision-making be informed by an accurate assessment of the risks of the particular transaction or customer concerned. Such models might be focused on fraud, using analysis of patterns revealed in past fraudulent transactions to predict how likely it is that this transaction is also fraudulent. They might be focused on the likelihood of default, using a customer's past payment history and the history of other customers like him to predict how likely it is that he will fail to make payments in a timely fashion.

Many techniques can be used to build such models from historical data, but all of them require knowledge of which historical transactions were "bad"—fraudulent or in default. These known cases are used to train a model to predict how similar a new transaction is to these "bads." Once such a prediction exists, a Decision Management System can use it, treating those transactions or customers with particularly high, or particularly low, risk differently.

Predict Opportunity

Many decisions do not involve an assessment of downside risk, but they still have some variability. Not driven entirely by compliance with regulations or policies, these decisions require an assessment of opportunity before an appropriate choice can be made. There is typically no absolute downside if a poor decision is made, simply a missed opportunity. When Decision Management Systems are being used to manage these opportunity-centric decisions, they will need to have some way to manage these tradeoffs.

These decisions are largely, though not exclusively, about how to treat customers. Deciding which offer to make to a customer or which ad to display to a visitor are examples of decisions where the "best" decision is one which makes the most of the opportunity to interact with the customer or visitor. Historical data can be used to predict how appealing a particular offer or product might be to a particular person or to a specific segment of customers. The value to the company of each offer, combined with the likelihood that a particular customer will accept it, can then be used to identify the most effective offer—to make the best decision.

When many such offers are being considered, it may be complex to identify the "best" offer. It may be difficult to manage the tradeoffs between the various

decisions. In these circumstances, Decision Management Systems can take advantage of optimization technology that allows the tradeoffs to be explicitly defined, and then the "optimal" or best outcome can be selected mathematically.

Predict Impact of Decisions

Sometimes the effect of an action taken by a Decision Management System cannot be precisely determined. For instance, the value of a subscription for a mobile phone will vary with the use made of the phone. When an action is available for a decision and has this kind of uncertainty about its value, a further prediction is needed.

The likely impact of each action on the profitability, risk or retention of a customer can be predicted by analyzing the behavior of other similar customers who were treated the same way—for whom the same action was taken. The prediction of the likely impact of each action can be combined with predictions of risk and opportunity to improve the quality of decision-making in Decision Management Systems.

Principle #4: Test, Learn, and Continuously Improve

Most information systems have a single approach to handling any decisions that have been embedded in them. Every transaction is treated the same way, with possible alternative approaches largely eliminated during design to find the "best" approach. Once this singular approach has been implemented, information systems continue to work the way they were originally designed until someone explicitly re-codes them to behave differently. The only way these systems are changed is when an external agent—a human—decides that a change is required. These systems also accumulate large amounts of data about customers, products, and other aspects of the business. This data might show that certain actions are more effective than others, but the system will continue with its programmed behavior regardless—every customer is treated like the first.

This approach is not an effective way to develop Decision Management Systems. When we make decisions about our own lives or interactions, we often assess a large amount of data, either explicitly or implicitly. We learn from this data what is likely to work or not

work—the data accumulated provides clues to how an effective decision can be made. A Decision Management System cannot afford to ignore the accumulated historical data.

Decisions involve making a selection from a range of alternative actions and then taking the selected action. It is often not immediately obvious if the decision was made effectively. Some decisions have a significant time to outcome, and no assessment of the effectiveness of the decision will be possible until that time has passed. For instance, an early intervention designed to ensure a customer renews her annual contract cannot be assessed until the customer reaches the renewal point, perhaps many months later. If the action taken turns out to be ineffective, then a different approach will need to be considered. A Decision Management System cannot afford to “single thread” this analysis by only testing one decision making approach at a time.

Whether a decision is a good one or a bad one is a moving target. A decision may be made to discount a particular order for a customer that may be competitive today but much less so tomorrow because a competitor has changed their pricing. As markets, competitors, and consumer behavior shift, they affect the effectiveness of a decision. This constant change in the definition of an effective decision means that Decision Management Systems must optimize their behavior over time, continuously refining and improving how they act.

Decision Management Systems must therefore test, learn, and continuously improve. The analysis and changes may be done by human observers of the Decision Management System or by the system itself in a more automated fashion. Decision Management Systems must collect data about the effectiveness of decision making. They must use this data, and other data collected by traditional information systems, to refine and improve their decision-making approach. Decision Management Systems must allow multiple potential decision-making approaches to be tried simultaneously. These are continually compared to see which ones work and which ones do not. Successful ones persist and evolve, unsuccessful ones are jettisoned. Finally, Decision Management Systems must be built on the basis that their behavior will change and improve over time. Decision Management Systems will not be perfect when implemented but will optimize themselves as time passes.

Collect and Use Information to Improve

The first way Decision Management Systems must learn is through collecting and then using information about the decisions they make. When a Decision Management System makes a decision, it should record what decision it made, as well as how and why it made the decision it did. This decision performance information will allow the long-term effectiveness of a decision to be assessed as it can be integrated with the organization's performance metrics to see which decisions result in which positive, or negative, performance outcomes. This information allows good decisions to be differentiated from bad ones, better ones from worse ones. It is often said that if you wish to improve something, you must first measure it. Decisions are not an exception to this rule.

Information about the decisions made can and should be combined with the information used to make the decision. This information might be about a customer, a product, a claim, or other transaction. This is the information that is passed to the Decision Management System so that it can make a decision. Combining this information with the decision performance information will identify differences in performance that are caused by differences in the information used to drive the decision. For instance, a decision-making approach may work well for customers with income below a certain level and poorly for those above it. Storing, integrating, analyzing, and using this data to improve decision-making is the first building block in building Decision Management Systems that continuously improve.

Support Experimentation (Test and Learn)

When a Decision Management System is being defined, it may not be clear what approach will result in the best outcomes for the organization. Several alternative approaches might all be valid candidates for "best approach." Simulation and modeling of these approaches, and testing them against historical data, might show which approach is most likely to be superior. Even if the historical data points to a clear winner, the approach is going to be used against new data and may not perform as well in these circumstances.

A Decision Management System, therefore, needs to be able to run experiments, choosing between multiple defined approaches for real transactions. The approach used for each transaction can be recorded, and this information will allow the approaches to be compared to

see which is superior. This comparison may not be definitive, and one approach may be better for some segments of a customer base, while a second works better for other segments. Results from these experiments can then be used to update the Decision Management System with the most successful approach or combination of approaches. Because Decision Management Systems handle repeatable decisions, there will always be more decisions to be made that will be able to take advantage of this improved approach.

Optimize Over Time

In a static world, one round of experimentation might be enough to find the best approach. A set of experiments could be conducted and the most effective approach selected. As long as nothing changes, this approach will continue to be most effective. However, the effectiveness of a decision-making approach can vary over time for many reasons, and you have little or no control over this. The old "best" approach may degrade suddenly or gradually, and when it does, you will need to have alternatives. Even when experimentation finds a clear winner, a Decision Management System needs to keep experimenting to see whether any of the alternative approaches have begun to outperform the previous winner. Alternative approaches could be those rejected as inferior initially or new ones developed specifically to see whether a new approach would be superior in the changing circumstances. The effect of this continuous and never-ending experimentation is to optimize results over time by continually refining and improving decision-making approaches.

Summary

Decision Management Systems are different from traditional information systems.

- v Traditional information systems have a process, data, or functional focus. Decision Management Systems are decision-centric, built with a repeatable decision in mind.
- v Traditional information systems are opaque and hard to change. Decision Management Systems improve collaboration and compliance by being transparent and agile.
- v Traditional information systems present historical data as analyses to people. Decision Management Systems embed analytics that predict risk, opportunity, and impact deep into the system

itself.

- v Traditional information systems are static and don't use the data they store to improve their results. Decision Management Systems test new approaches, learn what works, and continuously improve.

Developing Decision Management Systems requires a new approach; this is the subject of Part II, "Implementing Decision Management."