Quantellia Inc.
PO Box 36417
Denver, CO 80236
Tel: +1 720 242 7825
Fax: +1 775 295 9122

Decision Engineering White Paper

Overcoming the Decision Complexity Ceiling Through Design

Lorien Pratt and Mark Zangari
5 December 2008
Version 1.1

Modern decision making has become bound by a complexity ceiling

Organizational decision making has reached a *complexity ceiling:* the number of factors that come into play when making a major decision, along with their complexity, has exceeded our capacity to make the right choices[1]. The amount of information that must be considered, the number of choices faced, the way these can interact in unforeseen ways to produce disturbingly unexpected results, all accumulate to overwhelm the *ad hoc* techniques that most decision makers typically rely on. Couple this to the seriousness of the *consequences* that may accompany bad decisions—today often measured not only in dollars but lives—and it becomes clear that huge risks are being taken based on justifications that are often far from sound. The decisions made by large organizations, and society as a whole, have reached a level of complexity that has outgrown the capabilities of informal decision-making processes. The stakes have become too high and the game is now played too fast for relying on intuition and luck alone. We need a system that gives us the best chance of winning.

Data-driven analysis has been regarded as irreconcilable with knowledge based on experience.

The complexity that limits the effectiveness of decision makers comes in many guises and crosses many industries. One classic dilemma is the seemingly irreconcilable separation between data-driven, "analytic" approaches to supporting decisions, vs. "inductive", intuitive approaches. Both have their own strengths and limitations, but rarely are the two harnessed to cooperate within a unified decision-making framework. Generally, they compete, with justifications based on "gut reactions" regarded, at least from a methodological standpoint, as at odds with "rational" analysis. In fact, this dichotomy is not only artificial, but often, is detrimental to decision outcomes.

Examples of the challenges decision makers face along the lines described above are in plentiful supply. To name a few: after examining the principal sources of their revenues, telecommunications companies are re-evaluating what has been industry gospel regarding their position in the media value chain[2]; intelligence analysts recognize the need to change decision making strategies to respond to a War On Terror that is more organic and changeable than the Cold War[3]; financial markets have suffered by not recognizing the limitations of statistical models; physicians are moving to employ an *evidence based* paradigm in clinical practice[4]. This set is but a small sample.

Complexity permeates decision making from a number of sources. For example:

- The number of factors to be considered in the decision, including inputs, desired results, dependencies between elements of the decision, peripheral (often unintended) consequences, long cause-and-effect chains, etc.
- Time variation. Many of the above factors will change during the lifetime of the decision making and execution process.
- Data, which may be difficult to obtain and manage, difficult to interpret, only partially available, uncertain, or possibly simply incorrect.
- Human factors, such as differing viewpoints, levels of skills and experience of decision contributors, the effects of political and other social relationships, etc.

Common responses to decision-making challenges are clearly inadequate

Common responses to these challenges include: gathering more data, hiring experts (statisticians, modelers, and others), building more complex statistical models, or placing greater demands on IT—often in the form of building an information architecture that allows greater sharing and collaboration. Although these approaches may sometimes be helpful, they are not enough and fail to address the root of the problem. These solutions lack *integration* between the facts and figures used to make choices, human expert judgment, and the technology platform, if any. In many organizations, systems, data, and human stakeholders are separated by culture, language, geographical distance, and time delays.

Fortunately, as a technological society, we have faced this set of circumstances many times before, and a solution pattern has emerged that has proven to be very effective in

Following historically proven techniques, a methodical approach that fosters a synthesis of human judgment and data effectively overcomes complexity

many fields of human endeavor. Construction, software, finance, manufacturing and many other disciplines, as they advanced, all reached stages where the number of inputs, the interactions between them, and the timeframes in which results had to be achieved overwhelmed the techniques of the day. With that, the likelihood of successful outcomes, let alone further advancement, dwindled. In each of these cases, the complexity ceiling was overcome by a common response: the development and widespread adoption of a methodical process that:

- Systematized and made objective the tasks required for successful completion,
- Defined quality control checks and balances for each set of tasks,
- Defined a common formal nomenclature that removed ambiguity and facilitated information and knowledge sharing among groups with widely different skills and backgrounds.

Following the terminology from these other disciplines, we shall refer to this as an "engineering" approach, and propose that the development of decision engineering will provide the tools necessary for decision makers to break through the complexity ceiling. The Decision Engineering process includes[5]:

Decision Engineering includes:
- Decision Objective Statement
- Specification
- Iterative Design
- Change Control
- QA & Security
- Control and feedback from decision execution

- Recognition that the decision making process follows a *lifecycle* beginning with the formulation of the problem and ending with the completion of the solution.
- Clear definition of the decision objective and translation of these into formal specifications.
- An iterative design process that incorporates data and expert judgment, allows for multiple data scenarios, and models different potential worlds. Documentation deliverables produced by this step include Decision Maps, Models, and World Diagrams.
- Quality assurance and security steps at each stage.
- The definition of a powerful visual representation for all aspects of the decision making process, in particular, for facilitating the integration of expert knowledge with empirical data.

One of the key activities of the iterative design process is *sensitivity analysis*. As mentioned above, traditional "decision support" strategies typically attempt to fill information holes by the acquisition of yet more data, or the introduction of yet further model complexity. However, simply gathering data without first qualifying which data *most affects the decision* is a recipe for wasting resources by obtaining irrelevant data that, often despite strong intuitions to the contrary, will not actually inform the decision one way or another. By contrast, Decision Engineering measures the sensitivity of outcomes to the various inputs that affect the decision. In this way, data acquisition resources can be targeted to where data will provide genuine insight into the system modeled by the decision.

Furthermore, sensitivity analysis reveals the situations in which the decision is vulnerable either to changing circumstances, or to the fragility of conclusions that were reached with incomplete data. This provides decision makers with a number of tools, including a set of key leading indicators to monitor during execution of the decision, and a risk-assessment framework, identifying the conditions where the chosen world-model will not lead to the desired outcomes. In such cases, the framework specifies that alternative world-models also be developed that produce the desired outcomes under the circumstancces that are adverse to the initial model.

Visual software tools will enable Decision Engineering to yield maximum benefits.

Finally, Decision Engineering will deliver maximum benefits when used with a set of software tools that implements its methodology, automates many of the tasks, and provides a rich, visual application environment for decision makers and their collaborators. Quantellia's *World Modeler* software is the first such tool, providing comprehensive support to all phases of the Decision Engineering process.

Introduction

This white paper demonstrates how decision design overcomes the complexity ceiling within the broader Decision Engineering framework.

As an illustration, we describe here the approaches taken by two fictitious (but fact-based) telecommunications companies, which we call "AveryTel" and "ArwenCom". AveryTel uses a traditional approach. ArwenCom designs its decisions visually. We illustrate its approach by showing its decision models, maps, and world diagrams—the "blueprints" of decision design.

Although this example comes from telecommunications, it illustrates similar problems—and solutions—faced in many organizations, private and public, worldwide.

Note that we focus here on the decision design process. Other important aspects of decision engineering, such as the management of uncertainty, sensitivity analysis, quality assurance for decision models, decision model security, the role of linear and nonlinear predictive analytics, web-based collaborative review of decision models, and the proper role of judgment and intuition in decision making are covered in more detail by other papers in this White Paper series.

Table of Contents

The Decision: Should We Offer an Unlimited Monthly Subscription Service?

Consider a fictitious (but typical) mobile communications provider that charges by minute of talk time, and by kilobyte of data usage. The company, AveryTel, is unsure whether to launch an "unlimited" subscription service, where customers can use all the voice-call minutes and data they desire for a single monthly price. (Note that this example is a composite of several telecoms interviewed during a research study conducted during the early months of 2008. It reflects real choices made by several large companies.)

The president of the company convenes a team to analyze the costs and benefits of the proposed "Unlimited" plan.

A fictional company—named AveryTel—runs into problems when launching a new pricing plan

FEASIBILITY ANALYSIS

The AveryTel team begins by creating a spreadsheet to evaluate the plan costs and benefits. It also commissions a market research study to test consumer interest. The study does not include a question to determine the price that consumers would be willing to pay for such a plan, however.

The spreadsheet and market research demonstrate a strong consumer interest in the plan, so the AveryTel president then asks the team to roll it out throughout the company. This will impact sales, marketing, network operations, and many other functions.

AN OBJECTION FROM SALES

As the rollout moves forward, the sales department raises an objection: it is concerned that it may lose sales commissions under the new business model. However, rather than expressing his concerns regarding the effect the plan may have on his and his team's incomes, the vice president of sales instead decides to undermine the plan by building his own spreadsheet. Contrary to the project team's conclusions, the sales VP's analysis shows revenue losses in the millions of dollars to AveryTel should it continue to implement the Unlimited plan.

The company President intervenes, falling back on his expert judgment to cancel the project

COMPARING THE FUTURE

AveryTel's President evaluates the two spreadsheets to determine how to resolve this issue. He is quickly overwhelmed by thousands of cells to interpret, each with formulas and data reflecting different—and often buried—assumptions about how the new offering would affect the company, its competitors, and its customers.

Frustrated and overwhelmed, the President eventually abandons the spreadsheet analysis. Reasoning that he was hired for his expert judgment, he considers that a) such a business model has never before been successful, and b) a worldwide recession is imminent, so consumers will be hesitant to commit to a monthly purchase amount. Combining this with his belief that the cost of implementation will be considerable, he cancels the rollout of the Unlimited plan. A few months down the road, the economy improves, and one of AveryTel's competitors, ArwenCom Telecommunications, launches its own unlimited usage mobile telephone service. ArwenCom steals a considerable number of customers from AveryTel and, as the only provider of an Unlimited plan, exploits its position by reaping healthy margins from the product.

A second fictional company—ArwenCom—uses Decision Engineering to make the same decision

Decision Engineering Step 1: Determining the Decision Objective

We will now review how the same decision was made by ArwenCom, which used a Decision Engineering approach.

The first step in ArwenCom's decision engineering process is to determine the *Decision Objective Statement:* a written statement, in the language of the decision "customer", describing what the decision must achieve, as well as the latitude that decision makers have in choices to achieve that outcome. The decision objective constitutes the "marching orders" to the decision making team. ArwenCom's CEO—we'll call him Mr. Smith—created the team to evaluate the Unlimited plan, and so he played the role of this team's customer. Its mission was stated simply at first:

Determine if offering an Unlimited Usage plan would be a good idea for ArwenCom Telecommunications

Although this decision objective is vague, it is nonetheless typical of the initial request often given to decision-making teams by their customers.

During Decision Specification, the team works with its customer to determine exactly what it is to achieve, and what choices it can recommend to achieve it.

Decision Engineering Step 2: Specification Analysis

Since the Decision Objective Statement from Mr. Smith left out many of the details that the decision team needed to know, the team asked to meet with him for clarification. Although they were interested in Smith's ideas for *how* to implement the unlimited usage plan they explained at the start of the meeting that the primary focus of this meeting was *what* was to be achieved.

Nonetheless, Mr. Smith began by explaining that he thought that $50 per month might be a good price to charge for such a service. The team leader asked "are you requiring that this be what we charge, or are you simply suggesting that number as one option?" Smith answered that $50 just seemed like a good idea, based on his experience. In response, the team leader wrote on the whiteboard under the heading "Specification": "Determine the following aspects of an Unlimited Usage Plan". He then wrote one bullet, labeled "Price of the new service".

Decision Specification focuses on the "what", not the "how" of the decision.

On another part of the whiteboard, the team leader wrote: "Ideas for Decision Planning". The first bullet under this list was "$50 might be a good price". He explained to Smith that his judgment was very valuable, and that this price idea would be taken into account by the team as it weighed other evidence for the right price to charge. However, the team leader explained that it was important to keep separate the "good ideas" from the "rules of engagement" for the decision making team, so that the latitude the team had available to it was clear.

The ArwenCom team meeting continued as a brainstorming session. Decision team members were asked to suggest what "levers" they might have available to them that could affect their customer's interest in—and ultimately the success of—the new service. They came up with "Contract Period" for customers on this service, and "Inducements offered", such as a free handset. Mr. Smith agreed that he wished the team to determine the best values of these choices for ArwenCom

One team member suggested that the market to which the product was offered initially might make a difference to its success. He suggested that, if marketing dollars were restricted, that perhaps offering this plan to a European market first might produce better results. Smith explained that, although this might be true, there were a number of complications to offering this new service in Europe. As with the $50 pricing question, above, the team leader asked "is that just a good idea, or would you rather this team not explore options outside of the United States?". Mr. Smith answered "this is going to need to be part of the specification for this team. Please do not consider marketing outside of the United States." This aspect of the specification, along with the other choices available during decision making, was recorded on the whiteboard as shown in Figure 1.

Unlimited Usage Plan
Decision Specification

Determine the following aspects of a monthly unlimited mobile telephone service:

- Price of the new service
- Inducements offered (e.g. free handset)
- Contract Period, and
- Target demographic (limited to US markets) and marketing plan
- Amortization period for startup costs

Do not launch to markets outside of the United States.

Figure 1: ArwenCom Decision Specification, First Draft

A Decision Specification contains both choice points and desired decision outcomes.

The discussion then turned to the outcomes that the CEO, wished the team to achieve. Recall that the Decision Objective Statement initially given to the team by Smith simply wished to know whether this Unlimited Usage Plan was "a good idea". The team leader began this part of the discussion by telling Mr. Smith that he guessed that "good idea" meant that the company would achieve higher net revenues with the plan than without it. Smith then explained that he thought that a competitor might launch a similar plan and steal market share from ArwenCom. The team leader asked "would you like us to consider offering an Unlimited plan even if a competitor does not launch one of their own?" Smith's answer was that yes, ArwenCom needs to be *anticipating*, not just *reacting* to its environment, so the team should consider this scenario as well.

The team also discussed the meaning of "net revenues". The team leader asked Mr. Smith: "How will we know if we are successful?" Must we break-even on our startup costs at 12 months, 24 months, or some other time in the future? Smith was hesitant to name a specific time in the future. He said "I just want the ongoing benefit from this to exceed its cost, whatever time in the future we are discussing." The team leader explained that it was important that goals be clearly measurable. Smith continued to balk: "I just want this to be an ongoing success". The team took a short break, and on returning one member said "I think I know the answer here." What if we thought of our costs in two parts: the 'start-up' costs for this effort, which we will incur only once as we change our systems, marketing, sales, and other organizations to support this plan, and the 'ongoing' costs? I think that what Mr. Smith is saying is that the ongoing costs should be less than the ongoing revenues, even if for a while, this project runs at a loss while we incur start-up costs."

It is important to agree on definitions of important terms, as a starting point to decision design.

Everyone agreed that breaking up the costs in this way was the right choice, but chose to use the language "operational costs" instead of "ongoing". The team concluded with a discussion of the remaining goals to be achieved by the decision making process. Ultimately, the team and Mr. Smith agreed to the specification shown in Figure 2.

The team leader copied this specification from the whiteboard and sent it in a memo to Mr. Smith. In this memo, he included a section labeled "Glossary", where he defined each word in the specification that he thought might need clarification. For instance, he realized that the word "demographic" in the specification wasn't completely clear: did it refer to a geographic, age, or other segment of the population? He clarified his understanding in the specification memo to mean geographic targeting.

The completed Decision Specification includes multiple goals, multiple choice points, and measurement criteria.

Unlimited Usage Plan
Decision Specification

Determine the following aspects of a monthly unlimited mobile telephone service:

- Price of the new service
- Inducements offered (e.g. free handset)
- Contract Period,
- Target demographic (limited to US markets) and marketing plan,
- Amortization period for startup costs

Our goal is to identify the right values for these choices such that the net operational (ongoing) revenues exceed those that we will achieve without the plan, and also so that these net revenues are maximized. This should be true whether or not our competitor launches a similar plan. It may or may not be possible to identify such a plan.

We must achieve this goal while also continuing to serve other company goals. These are:

1) Maintaining employee morale
2) Long-term company viability—we will not 'rob Peter to pay Paul'),
3) Maintaining low environmental impact
4) Service to the community
5) Do not launch to markets outside of the United States.

We will measure our impact on these other company goals through an executive decision review process, where an executive committee checks our recommendations against these goals before decision implementation time. Our success criterion: unanimous agreement that our recommendations will not adversely affect these other goals.

Figure 2: Final Specification for ArwenCom Telephone's Unlimited Usage Plan Decision

As in other design arenas, in decision design a diverse team works in a structured process towards a common goal.

Decision Engineering Step 3: Decision Design

After reviewing the specification memo, the decision team felt that it clearly understood its mandate, and so began the *decision design* process.

Decision design recognizes that the creation of a decision shares many common elements with the creation of other engineered objects, like buildings, bridges, and software. To build any one of these objects requires that a number of specialists work together to reason about how best to combine basic elements (doors, wires, joists, walls) into something more complex. A decision is such a complex object.

Decision components can be represented visually.

Core to this approach: throughout its lifecycle, a decision is documented **visually** by representing basic components (like assumptions, external data, choices, outcomes) as a picture. Down the road, this visual depiction of the decision also serves as a basis for collaboration, organizational alignment, and ultimately, as a reusable object to facilitate a change in the decision based on a change in circumstances.

Decision *design* is the heart of the decision engineering process, and includes methodologies, tools, and processes for creating the decision itself. As with an architect's blueprint or software developer's UML diagram, a structured visual language for expressing decision elements, their properties, and the relationships between them, is critical. Such diagrams unambiguously both record and communicate assumptions, data sources, conclusions, and flow of reasoning.

The decision team creates a decision iteratively, through gradually adding detail over time.

Core to effective collaborative decision design is an *iterative refinement* process, where the design begins as a rough sketch to which detail is gradually added. For instance, a product's demand curve may initially be represented simply as demand being inversely proportional to price (with the proportionality constant unknown). As more data are acquired, or more analysis done, this can be refined and represented as a curve that can be sketched as a hand-drawn graph; finally, when the model is complete, these approximations may be replaced by a mathematical function that precisely describes the model's representation of the demand curve.

INITIAL DECISION MODELING

The decision design team worked with an interactive electronic whiteboard. The team leader began by adding boxes representing each of the choices available to the team on the left-hand side of the screen. These were copied directly from the specification, and appeared as shown in Figure 3.

(Note that this and subsequent whiteboard pictures were created by a pre-alpha release version of Quantellia's WorldModeler™ tool. As this tool matures, we envision a greater degree of visual sophistication, including pictorial icons, and static and animated displays of high-dimensional information, as appropriate.)

The first step in decision modeling is to draw decision choices and outcomes

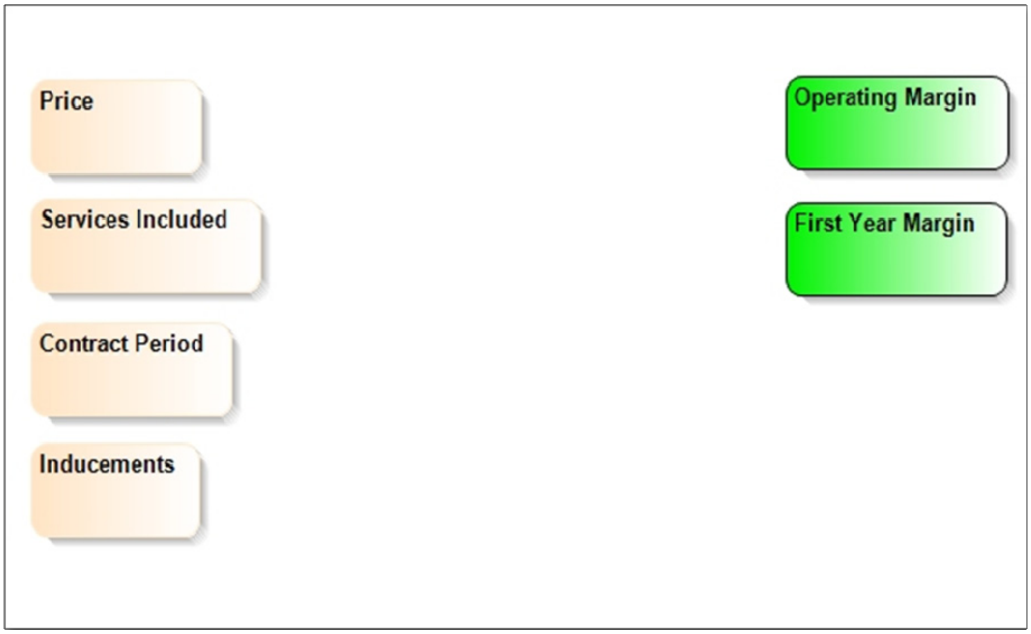

Figure 3: Decision Team's First Drawing: Decision Modeling Choices and Results (produced by WorldModeler software)

Thinking through the most recent discussion with the CEO Mr. Smith, one team member stood up to add a few intermediate values to the diagram, representing the revenues and costs that would lead to the margins. He drew lines – called *dependencies* representing which outcomes would be affected by those intermediate values. After rearranging a few elements to make room, the white board appeared as shown in Figure 4.

Intermediate values begin to show how choices lead to outcomes.

Dependency lines show information flow through the decision.

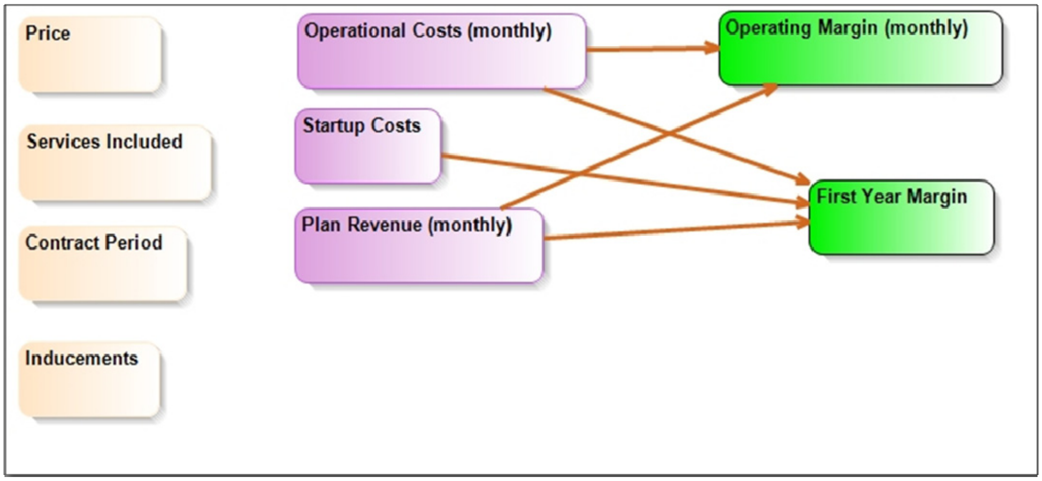

Figure 4: Enhancing the Decision Model: Adding Intermediate Values (produced by WorldModeler software)

At this point, the ArwenCom team realized that this model was too complex to show on a single diagram. It also realized that the diagram did not capture all elements of the specification – it left out the marketing and the amortization period choices. The team decided to take two steps to break down this problem into manageable chunks, as described in more detail below:

1) It would break the single decision down into several decisions that would be made over a period of time.

2) It would decompose the decision model based on outcomes

One way to reduce diagram complexity is to break one decision into a sequence of decisions over time, represented on a *decision map.*

BUILDING A DECISION MAP

To accomplish (1), the team drew a new kind of diagram – called a *decision map,* shown in Figure 5. A decision map is a number of ovals—each representing a decision—connected by lines each of which represent a temporal or logical dependency. This map shows that the team chose to separate selecting the attributes of the unlimited price plan product from the decision as to how to market it, and from the decision as to how to amortize the startup costs, and that they will make these decisions in this order.

Each decision on a decision map is defined by the choices that are made as part of that decision. These are shown in small rectangles in Figure 5.

The team realized that the marketing decision would probably result in a number of choices that were not represented on the decision map. It chose to defer considering marketing in depth so that it could focus on decision modeling for the "Define Product" decision.

DECOMPOSING A DECISION MODEL

As mentioned in (2) above, "Define Product"—the decision started in Figure 4—is too complex to represent on a single diagram. The team chose to focus at first on just the elements of the decision leading to calculating the startup costs for the project. Again, they drew the decision choices on the left and the outcome—in this case startup costs—on the right. Figure 6 shows the result.

A *decision map* shows how a sequence of decisions are related over time

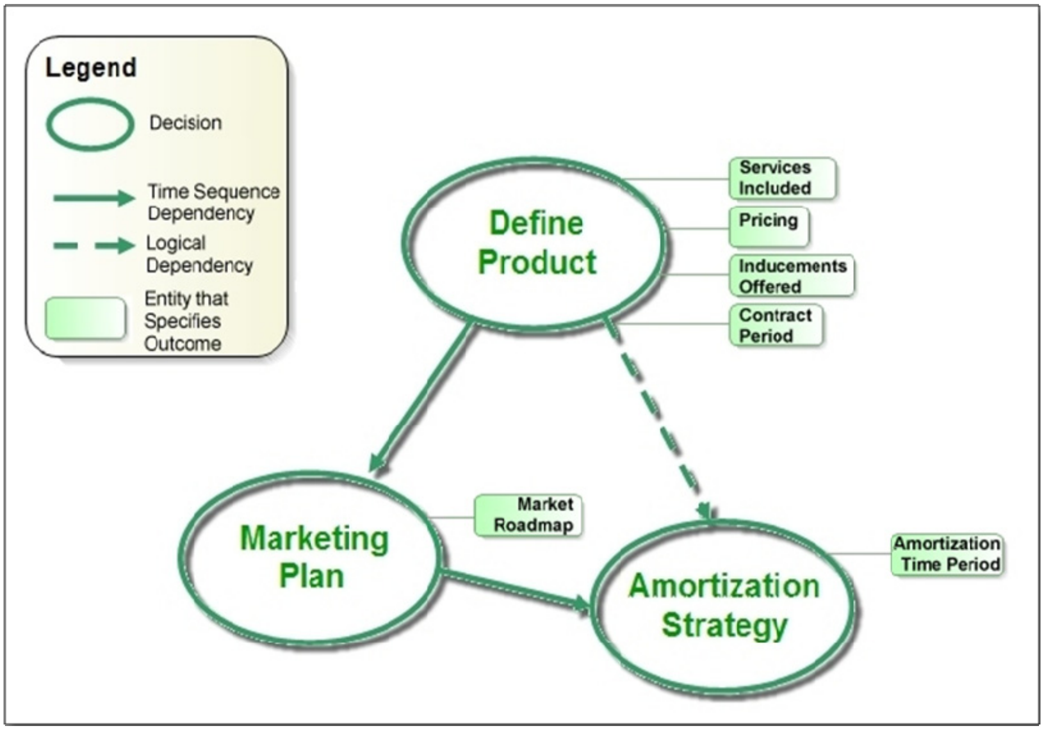

Figure 5: Decision Map for Unlimited Pricing Plan Comparison (produced by WorldModeler software)

A decomposed decision model might contain only a single outcome

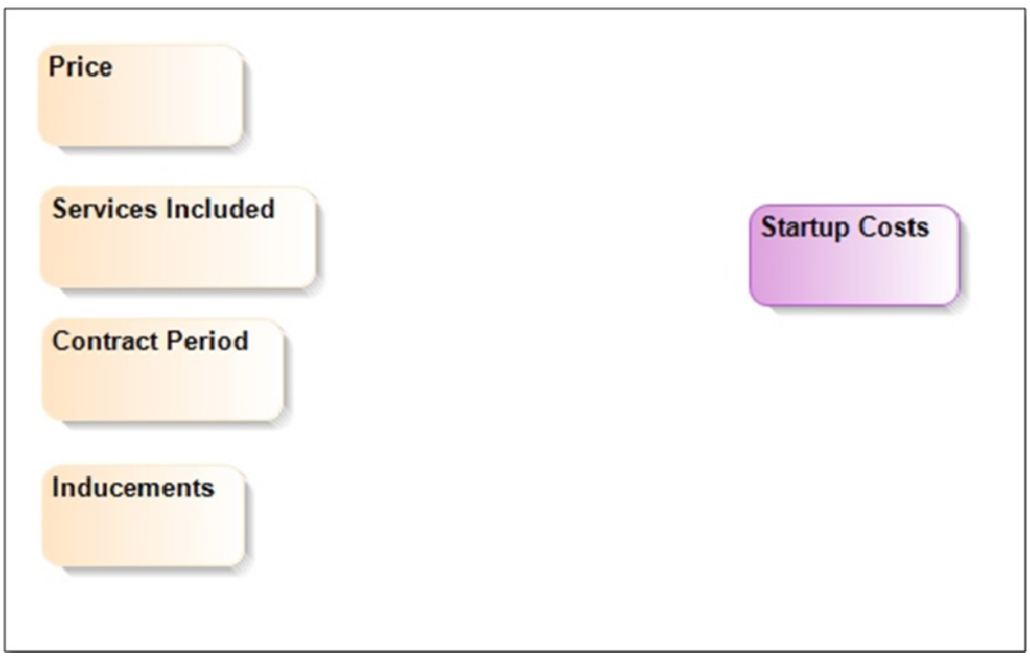

Figure 6: Startup Costs Sub-Decision Model: Choices and Outcome (produced by WorldModeler software)

The team next considered what costs would be involved in starting up the Unlimited service. They realized that changes would pervade the company, in many departments. They represented those departments as shown in Figure 7. The small orange boxes labeled "+" on this map indicate *composition:* that startup costs have four parts: operational, sales, customer care, and marketing startup costs. The cost of marketing changes consists, in turn, of two elements: the cost of a one-time marketing campaign to announce the new product, plus one-time costs to change the marketing department so that it can support this product on an ongoing basis.

The team decomposed "startup costs" into the components of those costs.

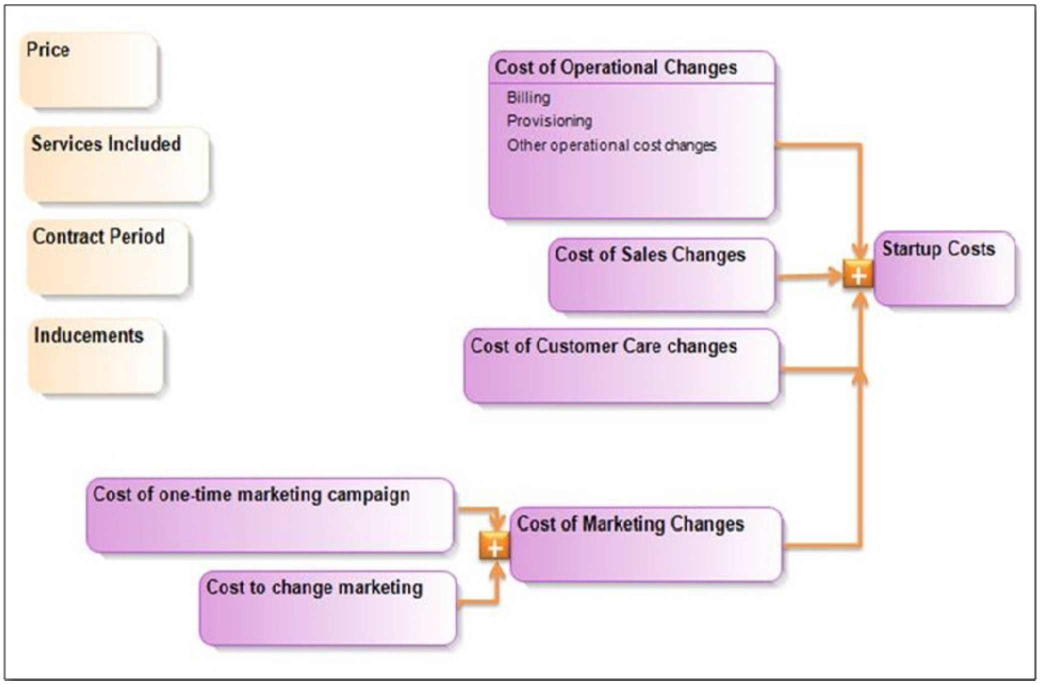

Figure 7: Startup Costs Sub-Decision Model, showing dependencies (produced by WorldModeler software)

Arrows representing dependencies show which values are determined by which choices.

The next step was to determine which of the decision choices might impact some aspect of startup costs. Upon reflection, the team realized that startup costs would be the same no matter what the price of the product, the contract period, or the inducements offered. The only choice that might impact startup costs was the services included in the product, because if there were a lot of services included, this would mean a bigger change to the operations systems like billing and provisioning, to sales, and to customer care. The team decided that the same change to marketing would be needed if there were one or a dozen included services in the plan, so determined no dependency here.

They represented these dependencies as the blue lines shown in Figure 8.

Before identifying the complete formula relating one value to the other, the team agreed to the *direction* of force between model elements

ADDING DIRECTION-OF-FORCE ARROWS

Although to understand this model fully, the team would ultimately need to place some sort of formula on each connector between entity boxes, the team began by simply determining the *direction* of impact (that is, whether the dependencies should be approximated as proportional, or as inversely so). If, for instance, the greater the number of services included in the unlimited plan contract, the greater the cost of sales changes (i.e. a proportional dependency), then the dependency was annotated to indicate this, with an arrow containing two plus signs read as "plus leads to plus". If, on the other hand, a rise in one value led to a *decrease* in another (an inversely proportional relationship), the annotation was an arrow containing a plus and a minus sign read as "plus leads to minus".

Since the cost of a one-time marketing campaign and the cost to change the marketing department to handle the unlimited plan did not depend on any choices made by the team, another change was to change the color of these entities in the decision diagram. The diagram now contained three kinds of entities: choices (shown on the left), Direction of force can be represented as "+..+", indicating a positive impact, or "+..-", indicating negative impact.

intermediate values, and these new "Assumption" entities, reflecting assumptions about an external value not under control of the team.

After completing these changes to the startup costs model, the resulting diagram is shown in Figure 8.

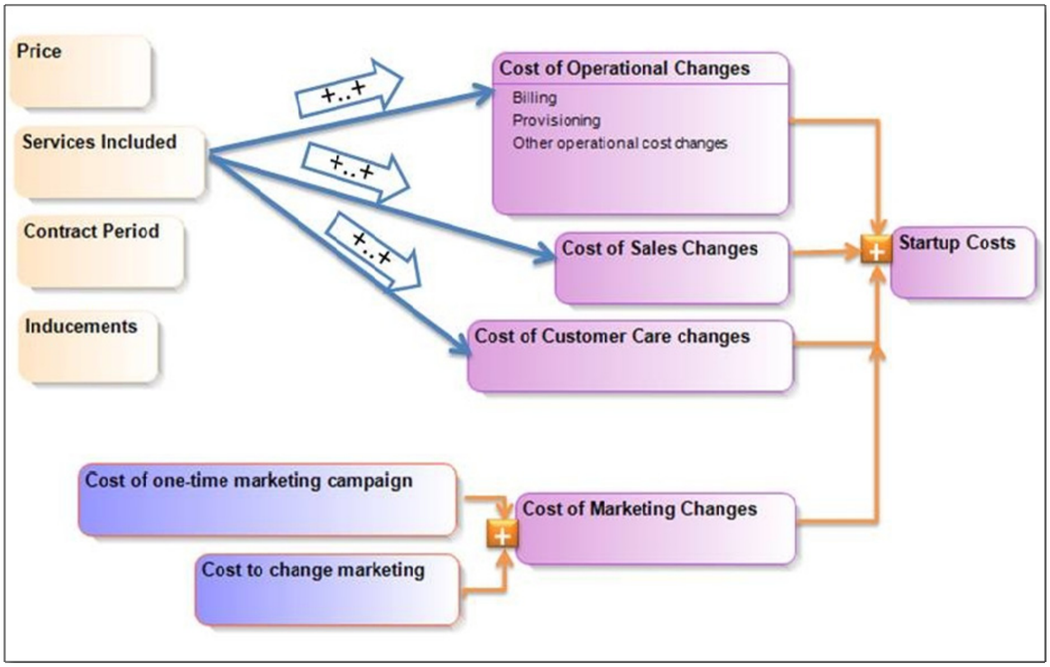

Figure 8: Startup Costs Diagram, including Dependency Direction Annotations (produced by WorldModeler software)

The team went through a similar process to build a decision model diagram for the ongoing network costs. The team's first version is shown in Figure 9.

Note that, in this model, operational costs consist of customer care, billing, network, and sales costs. Billing costs consist, in turn, of usage costs ("usage" means the collection of data records indicating the length and other parameters of a telephone call or the amount of data used by the handset), invoice costs, and other billing costs (this is a catch-all for functions like billing auditing).

Note the red arrow on this diagram pointing towards "Billing Question Costs". The team used this arrow to indicate their belief that more customers who purchase the unlimited plan, the fewer billing-related questions are expected to come into customer care (since disputing minutes is no longer necessary). One team member raised an objection, however. He said "certainly there will be *some"* billing related questions, where are those captured?" Thinking further, this team member became more confused. "Certainly the most customers we have, the *greater* the number of billing questions, not fewer, as this arrow shows."

BUILDING A DECISION WORLD DIAGRAM FOR MODEL COMPARISON

These questions puzzled the team leader for a moment, until he realized the team's mistake. What they really wanted to know was the *difference* between two different courses of action: one in which ArwenCom Telecommunications launches the unlimited minutes plan, and the other in which that plan is not launched. Thinking further, he realized that the situation would be different based on whether a competitor also launched such a plan, since the company could expect to lose some customers to the competitor, and there would be downward pressure on prices. All in all, there are actually four different situations that need to be compared, and it is the *comparison* between these situations, in addition to the dynamics of the situations themselves, that is of interest.

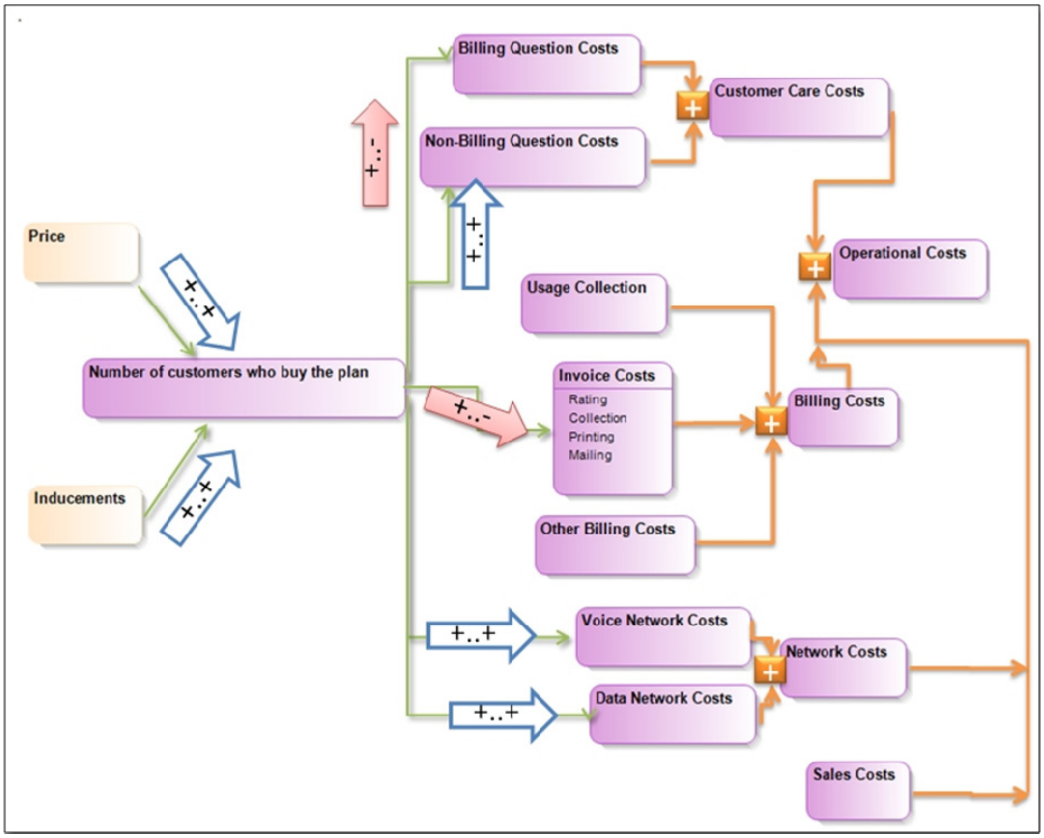

Figure 9: First Attempt at Operational Cost Decision Model (produced by WorldModeler software)

Based on this analysis, the team leader opened up a new window on the interactive whiteboard and drew the *decision world diagram* shown in Figure 10. In a decision world diagram, each ball-shaped "world" represents a different decision model. Differences from one world to the next may be between the structure of how entities are connected or in the data that feeds them. The diagram is drawn in an upside-down tree shape to show that different worlds share significant structure from one to the next.

For ArwenCom this means that, when the product evaluation team wants to compare its net revenues with and without the unlimited pricing plan, the company does not have to build a completely different model. Instead, there is a global "world" that contains most decision model elements, called "Telecom Product Model". Its "child worlds" have small changes that reflect the result of offering or not offering the Unlimited plan (red and green spheres in Figure 10). At the third level down, this figure introduces a new set of comparisons—not previously discussed—between how revenues play out if there is a competitor offering an unlimited plan and if no competitor offers such a plan. In this document, we leave the introduction of competitor effects as an exercise for the reader.

Now that the ArwenCom team understood that its primary goal was to compare the behavior of two different models, it corrected the diagram showed in Figure 10 by modifying the operational cost model in the red world (see Figure 10) to reflect costs associated with the unlimited pricing plan, and the model in the green world to represent the costs associated with providing their existing products. Both models now included new entities reflecting the number of customers on various plans.

A *world diagram* represents different—but related—model structures to be compared.

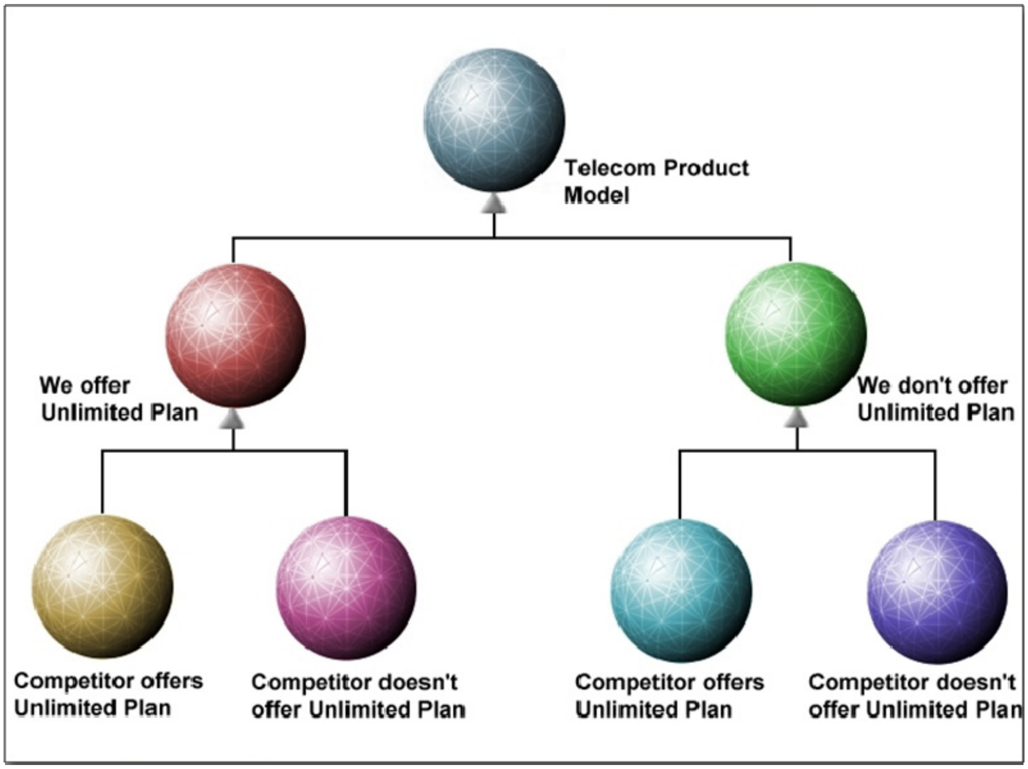

Figure 10: World Diagram for Unlimited Pricing Plan Comparison (produced by WorldModeler software)

To further refine the decision model, the team replaced force direction arrows with sketch graphs

CONVERTING FORCE ARROWS TO SKETCH GRAPHS

Examining these two models, the team realized that, at this level of detail, they were almost identical, and didn't really capture the key distinctions between operational costs with and without the unlimited plan on offer. After some discussion, they realized that the two red arrows that were shown in Figure 9 needed some refinement. Since customer care was a significant cost for this company (and one which was growing as competition increased), and since most customer care calls involved billing questions, ArwenCom realized that there were significant cost savings possible with the unlimited plan, where usage information would not be an issue, since it didn't impact a customer's bill.

Another important cost saving with the Unlimited plan was that processing usage information—which could be as high as a billion records *per day*—was less costly. Although the need to collect usage data remained for legal and interconnect billing reasons, the time requirements on usage processing as this information was collected and printed in a bill would be significantly reduced with the unlimited plan.

To capture these important distinctions, the simple "+..+" or "+..+" annotation was not adequate. Although the ArwenCom team felt it had more information available, it wasn't sure of the exact formula connecting the number of customers to billing questions and the number of customers to invoice costs. The team leader explained the use of a *sketch graph* in this situation, and showed the team how to replace the arrows on the decision models with graphs representing the approximate relationship.

Adding these elements created a new level of complexity, so the team decomposed a second time. It chose to focus for now on just customer care costs. The resulting model, which required about two hours for the team to build, is shown in Figure 11. This model looks complex, but each element is easy to understand when analyzed one at a time.

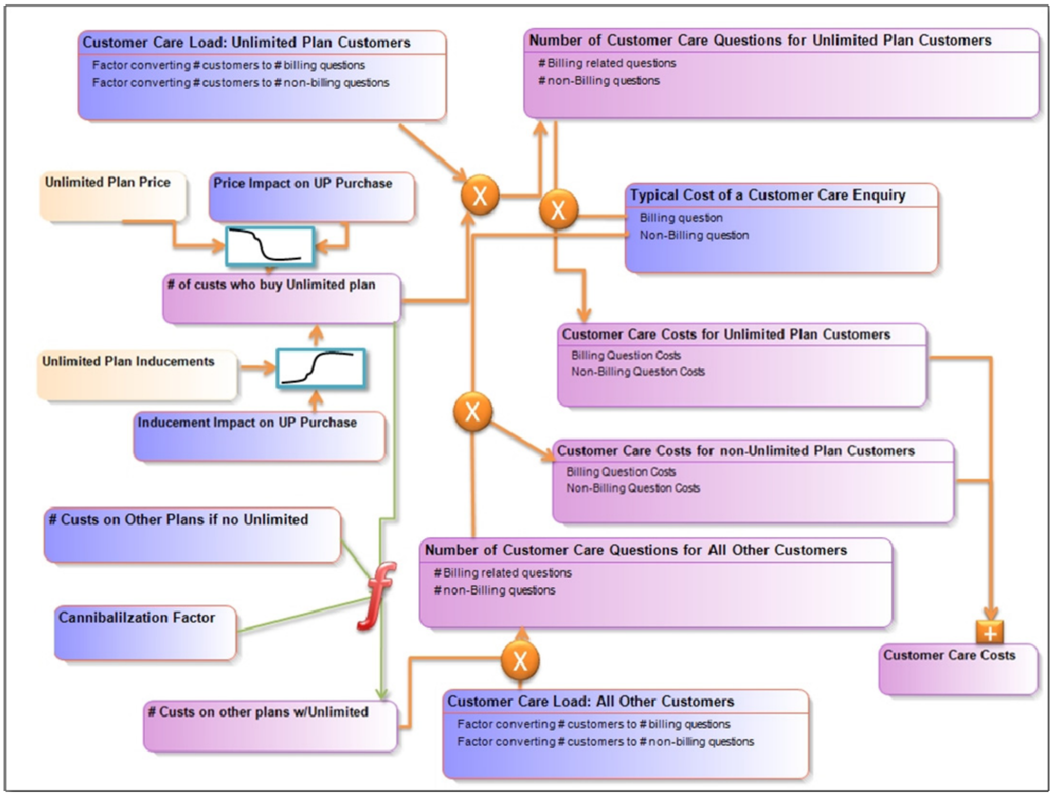

Figure 11: Customer Care Model, Showing Formulas and Sketch Graphs (produced by WorldModeler software)

> To create a sketch graph, the team drew the shape of the functional relationship between two decision entities

SKETCH GRAPH EXAMPLES

First, consider the green rectangle containing a black curve that connects **Unlimited Plan Price** and **Price Impact on UP Purchase** to **# of custs who buy Unlimited plan.** This graph reflects the fact that the team is not sure how users will respond to price as they consider the Unlimited plan. The graph shows, roughly, that there will be a large number of users interested in the plan if it is priced below some level, and that the number of users will drop to near-zero if the price is too high. The shape of this graph is determined by the factor **Price Impact on UP Purchase** as well as the shape of the function that relates this number to the number of purchasers. A similar sketch graph relates the number of inducements (free handsets, etc.) offered to customers to the number of customers interested.

> In addition to addition and dependency connections, a decision model also contains multiplication connections

MULTIPLICATION CONNECTORS

The orange circles containing the letter "X" represent a simple multiplication. The four multiplications in this model calculate:

1. Converting the total number of customers who purchase the Unlimited plan to the number of customer care questions that they ask, split into billing-related and non-billing-related questions
2. Similarly, converting the number of customers who purchase other plans to the number of billing and non-billing-related questions
3. Converting the number of questions asked by Unlimited Plan users to the cost of those questions to ArwenCom
4. Converting the number of questions asked by non-Unlimited Plan users to the cost of those questions to ArwenCom

A Function Connector represents how numbers are combined through a numerical function

FUNCTION CONNECTOR

Finally, this model also shows a red ***f***, which represents a more complex function than is represented by a simple multiplication, addition, or sketch graph. Although we have not shown this function here, the ArwenCom team designed a function that determines the number of customers who purchase other plans if an Unlimited Plan is offered. This is based on the number of customers on other plans today, a Cannibalization Factor representing how many of those customers will switch to the new plan, and the number of customers who choose to purchase the Unlimited plan.

External Values are numbers that the team cannot control.

CHOICES VS. EXTERNAL VALUES

The entities labeled **Unlimited Plan Price** and **Unlimited Plan Inducements** represent values that the decision team can select. Recall from the Specification phase that the ArwenCom CEO directed the team to determine the right values of these numbers.

In contrast, the blue boxes in Figure 11 represent data that, although also input to the model, we cannot control. These are called *external values.* External values may or may not be known at the time of modeling. If they are not known with certainty, then it is important to represent their *data quality*, as is discussed below in more detail below.

The team can explore multiple levels of data model abstraction by rolling up and drilling down on the model structure

DRILL-DOWN AND ROLL-UP

Although the decision model shown in Figure 11 appears complex, note that the orange symbols labeled "+" and "X" represent *roll-up points.* A tool to support decision design can be designed to allow the user to click on these symbols, which hides some entities on the screen, thereby exposing higher level model structure. This facilitates model understanding, especially at this stage of design, when many details—which may otherwise become overwhelming—have been added. An example is shown in Figure 12.

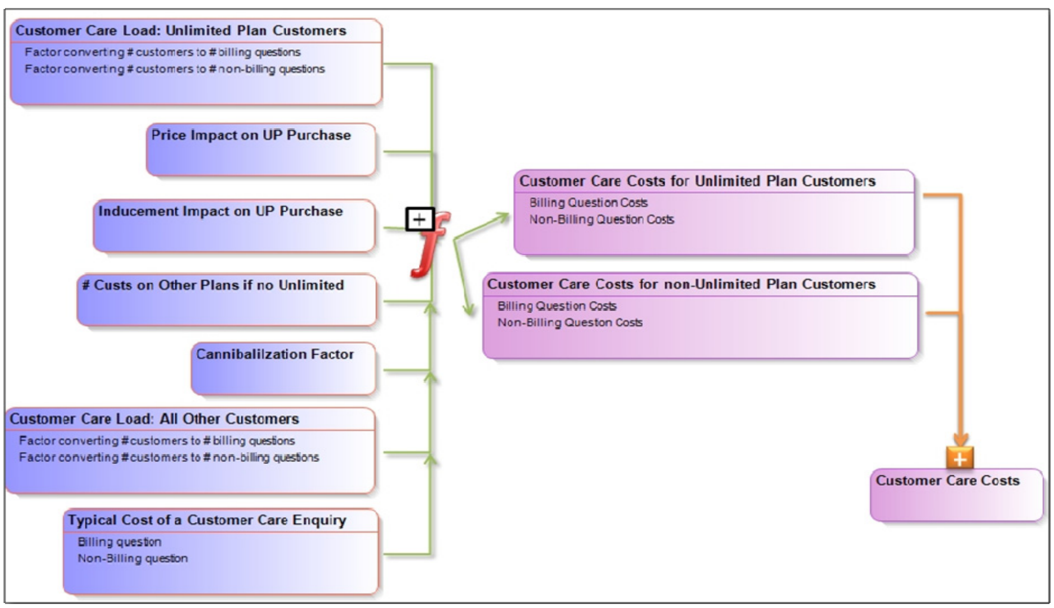

Figure 12: Rolled Up Customer Care Costs Sub-Model (produced by WorldModeler software)

BINDING DATA

Even at this stage, the modeling exercise has already yielded considerable insight to the team, regardless of whether it was performed on paper, on a whiteboard, or using a sophisticated software tools.

However, until now, we have been dealing with the abstract structure of the model. No values of any kind, numerical or otherwise, have been considered. However, ultimately, modeling (in this case) is about numbers, and very specific numbers at that, namely, the differences in revenues between offering an Unlimited plan and business as usual. How do we turn the abstract representation of the model in Figure 12 into numerical calculations?

Data Binding associates data values to model entities

The answer is *data binding*. Each of the blue entities in Figure 12 represents one or more data values that the team needs to plug into the model. It performs this "plugging in" by associating each entity with a *data source*. If the entity has *properties*, as do **Typical Cost of a Customer Care Enquiry** or **Customer Care Load** from Figure 12, then each property is associated with its own value from a data source.

Data binding is a key feature that gives the decision design approach much of its power and flexibility, in that:

- The model structure is always distinct from the data, making it easy to both view and manipulate the structure without affecting the data, and vice-versa. We say this makes the model structure *transparent* (as compared to spreadsheets, where data and model are combined in a single representation and, for the most part, the model structure is hidden or *opaque*).

Separating model from data provides a number of benefits to a decision modeling team

- There is a clear distionction between changing input data values, and changing the structure of the model itself.
- The same model can be fed from many different data sources. *Scenarios* may be established where a model entity is bound to any of the following types of data source:
  - Values entered manually into a table
  - A spreadsheet
  - A relational database query.
  - The return values of a web service
  - Custom data srouces developed using APIs

It is important to model *Data Quality* (including uncertainty, age of data, and other factors) because most decisions are made without perfect information.

  - The data binding information can also track metadata about the *quality*, *reliability* of and *confidence* in the data being mapped to the model structure. The modeling software can therefore assign a data quality metric to each entity in the model by using formulas to process not only the values of the bound data, but also the data quality metrics. The modeling software can then not only display the *value* of calculated quantities in the model, but also a measure of the *confidence* in each value, based on the quality of the data that was used to calculated it (note that management and propagation of data quality is a detailed subject in itself that lies outside the scope of this paper. See endnote 6).
  - Data bindings may be *static* or *dynamic.* Static bindings store a snapshot of the data from the source that, once populated, does not change. Dynamic bindings store only the specification of how the data is to be obtained (e.g. a range from a particular spreadsheet, or a particualr SQL query to a named data source), and refresh these values each time the model is run.

Based on the above, the ArwenCom team sought to find appropriate data sources to which to bind the input entities on the left hand side of Figure 12. The team found a summary spreadsheet produced by the customer care department that grouped the proportion of calls received each month into a number of categories. These were:

  - Billing related calls
  - Changes to the customer's product (e.g. plan changes, services added or removed)

- Subscriber address changes
- Technical Support

Fields in data sources may not map directly to the decision model. Data bindings create a mapping between the two structures.

As can be seen, these do not map directly onto the properties the team identified in the model (these being **Factor converting # of customer to # billing questions** and **Factor converting # of customer to # non-billing questions**). However, they can be combined in a way that gives the required results. The intention behind these two properties is to separate calls that are directly related to customers having separately billed product features from calls that are independent of the way the services are billed. The **Billing related calls** and **Changes to customer's product** fall into the latter category—these reduce dramatically if there is one and only one product—while **Subscriber address changes** and **technical support** fall into the non-billing related questions category, as it would expected that these will continue at the same rate regardless of how the products are structured.

As a result, the data bindings for the **Customer Care Load: All Other Customers** model entity is as shown in Figure 13, below.

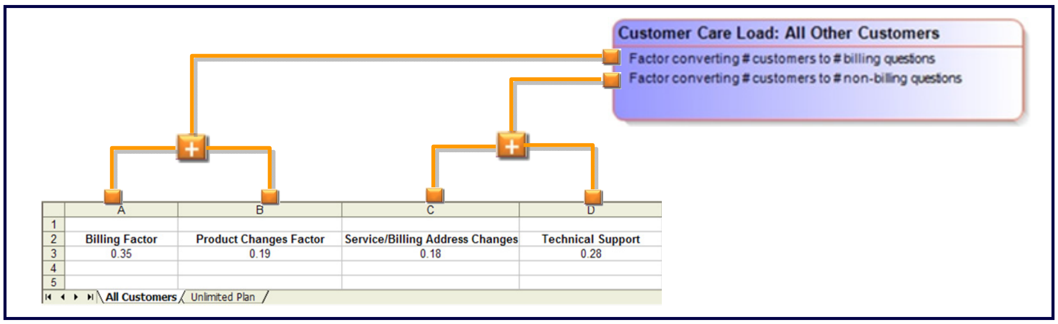

Figure 13: Binding Model Entity to a Data Source (produced by WorldModeler software)

Fields in data sources may not map directly to the decision model. Data bindings create a mapping between the two structures.

APPROXIMATING DATA VALUES / SENSITIVITY ANALYSIS

Obtaining exact numbers for the external values in a model is often very expensive, and so it is important to direct the majority of the data acquisition effort towards obtaining the data that has some impact on the final decision. Data approximation is a key point of differentiation between decision engineering and many other approaches. Here is the procedure:

1) Build a model, as in Figure 12, to the point where data elements needed for calculations are clear (they are shown in blue).

2) When an external value is acquired, bind it to the model in the approapriate place.

3) For each unknown external value, working with the team, determine the 90% confidence interval: within what range is everyone sure that the value will land? (Note that psychological research[7] indicates that humans tend to create intervals that are too narrow, but a collaborative approach can help here, especially if team members develop their estimates without consultation with each other, and then compare their estimates, taking the broadest uncertainty range so developed). This range can then be bound to the relevant model elements (again, with metadata indicating this value is an approximation, whose origin and history should be well understood, documented, and easy to track. See endnote 6))

4) **Run the Model**: Even if only ranges are known for each data value, and even if detailed formulas are not available for every dependency, the model can be run over

Sensitivity Analysis focuses data collection effort on the information that is needed to make a particular decision. Since data collection and maintenance can be very costly, this can produce a significant benefit, both in cost savings, and also focus on collection of the data elements that do have a large impact on the outcome of the decision.

a number of values in each data range to understand how the model outcome varies with different values of unknowns, and with different values of the model choices.

5) **Sensitivity Analysis:** Run the model to determine which unknown values have an impact on the model output, and which do not (it is here that a software tool becomes essential, to allow users to experiment with various parameters, and to automate much of this exploration process). For instance, it may turn out in the model shown in Figure 11 that there are so few billing questions generated by Unlimited Plan users, that the exact value of the cost of each billing question for an Unlimited Plan user does not need to be determined. The team may estimate that this cost is between $3 and $10, and the model may show that no matter where the actual cost lands within this range, it will not have a substantial impact on the customer care cost.

The value of being able to ignore collecting certain data elements is a great benefit of this approach, and should not be underestimated, and so bears repeating here. Data collection and maintenance in support of decisions is a very costly undertaking, and the ability to justifiably avoid these costs because refining an uncertainty range would have no impact on a decision has great value.

Note that, for simplicity, this example model is only a sub-model of the overall system under study, and we have only provided the flavor of the sensitivity analysis approach. If we were to fully develop this model, the criterion for whether a data value mattered or not would be whether it swung the decision to offer the plan one way or the other.

Key Assumptions are those external values that are both uncertain but that also have a large impact on the model outcome.

KEY ASSUMPTIONS

In addition to determining for which elements significant investment in better data is worthwhile, sensitivity analysis also indicates upon which values the model critically depends. These are sometimes called "Key Assumptions".[8] They are the most important values to obtain additional information about, if possible, during model building. Following the decision planning stage, when the decided plan is being implemented, Key Assumptions that are the most important elements to monitor as they provide forewarning that the likelihood of successful outcomes may be decreasing if the values tracked by the Key Assumptions change.

The decision model is not a quantitative modeling tool, but a team "intuition pump", aiding in alignment and coordination around difficult decisions that are typically fraught with uncertainty and nonlinear complexity, and so require expert judgment.

EXPERIMENTATION, INTUITION, CHOICES, AND ASSUMPTIONS

It is important to emphasize that the purpose of decision design is *not* to determine the optimal value for the model. This is, in most real cases, simply not possible: typical systems are very highly nonlinear and a finitely achievable optimization solution does not exist. Decision models should *not* be regarded as automated machines which, when fed with the right input data, will spit out the right answer. Rather, they incorporate such insights as are available by examining the data, and combine this with an environment where insights from experience can be easily be combined with the quantitative elements. For this reason, the judgment of expert decision makers is an essential ingredient in decision design.

Furthermore, it is not possible for expert decision makers' judgment to be easily extracted in a single step and handed off to a quantitative analysis team for input to a statistical model. Instead, to best leverage expert judgment, the expert team must work deeply, and interactively, with a representation of the decision, data, and facts that are available to them. The purpose of the decision design exercise is to support and align the team's intuitions about the decision at hand.

At this point in decision modeling, the team will begin to experiment with different model inputs to understand how they impact the outcome—ongoing revenues in this case—that they are trying to achieve. As the team experiments with the model, it may choose to:

Adjust values the organization can control (price and inducements in our example) to understand their impact on the outcome. A typical model design tool will offer interactive widgets to support this interactive exploration.

- Change the model structure. That is, modify the elements included in the model, their properties, and the relationships between them.
- Gather new data to narrow confidence intervals for external values. This may be done in response to sensitivity analysis that identifies key assumptions.
- Generally speaking, the team works iteratively and experimentally at this point, using the structure of the problem at hand rather than any strictly proscribed process as a guideline.

The ArwenCom team built models for the Unlimited plan, like the one shown in Figure 11, representing all components of the cost and revenue. The models included costs of network, operational, and sales (as were shown in Figure 9), as well as operational revenues.

COMPARING WORLD MODELS

The decision model is not a quantitative modeling tool, but a team "intuition pump", aiding in alignment and coordination around difficult decisions that are typically fraught with uncertainty and nonlinear complexity, and so require expert judgment.

As was shown in Figure 10, this team's goal is to compare a number of decision models which differ based on different model structures and/or different data. The key comparison between models that is of interest to ArwenCom is the different impact on revenues based on whether the company does or does not offer an Unlimited plan. It is interested in this comparison in the situation where a competitor does and does not offer a plan, as well. A decision design tool should allow multiple models to be displayed at the same time, and to offer facilities to detect and highlight model differences.

The ArwenCom team created several sub-models, each differing from the "parent" model to reflect the situation where an Unlimited Plan is and is not offered. It determined that, even if a recession occurred, demand for the Unlimited plan would be strong enough to warrant its cost. An unexpected outcome from ArwenCom's analysis was that adopting the Unlimited plan would result in an enormous reduction in customer care costs, via greatly reduced calls from its subscribers regarding billing.

ArwenCom saved significant market research costs, since it learned from decision modeling that savings in customer care for the Unlimited Plan swamped the value of any particular price for the Unlimited plan.

During the modeling exercise, the proportion of customer care calls related to billing was identified as a parameter to which the decision was sensitive, so the team carried out an extensive data mining and analysis project to make sure their numbers were sound. During this process, the team discovered that the lion's share of customer care costs came from usage related billing disputes. Since customer care was a labor-intensive—and therefore expensive—part of ArwenCom's operations, removing this cost component turned out to significantly "lower the bar" for success for the Unlimited plan. In fact, unlike AveryTel, ArwenCom did not perform significant market research since the effect of variations in consumer sentiment was swamped by the call center cost savings.

In a separate modeling exercise, ArwenCom chose whether to re-assign or lay off the customer care employees no longer needed for billing questions. It found that long-term relationship building was a worthwhile investment with ArwenCom's new crop of high-value corporate customers, and so was able to justify re-deploying the customer care professionals, along with a retraining program.

PRESENTING CONCLUSIONS

Once ArwenCom's team was convinced that the Unlimited Plan would be a success, it returned to Mr. Smith. Instead of a traditional PowerPoint, Mr. Smith was invited to the model building room, and the team leader walked him through the model structure on an interactive whiteboard, and then showed how the team experimented with different model values to reach its conclusion.

An important purpose of a Decision Model is as a mechanism for explaining the rationale for a decision to stakeholders.

Mr. Smith disagreed with the team's assumptions regarding an upcoming recession, and in response to his input, the team modified this external value. They showed him, however, that even under the most conservative of assumptions, including a slowdown in the economy, vigorous competitor responses, and slow adoption, the cost savings from the Unlimited plan still justified its adoption.

COMPANY-WIDE REVIEW

The team then posted the model on ArwenCom's intranet for employee commentary. Several customer care managers were able to identify a mistaken assumption in customer care costs for non-billing-related questions. Another employee had heard a rumor about a competitor's Unlimited plan, so she was able to make an adjustment on the Intranet model and submit it for approval to the assessment team, where the change was accepted. A third employee identified an additional low-cost inducement that the team agreed would increase net revenues by 3% for very little outlay.

An important purpose of a Decision Model is as a mechanism for explaining the rationale for a decision to stakeholders.

After the company-wide review, a new team was convened to roll out the Unlimited plan to ArwenCom as a whole. It was met with broad acceptance.

Additional Decision Engineering Steps

This concludes our overview of Decision Design, which is one part of an overall Decision Engineering methodology, as illustrated in Figure 14.

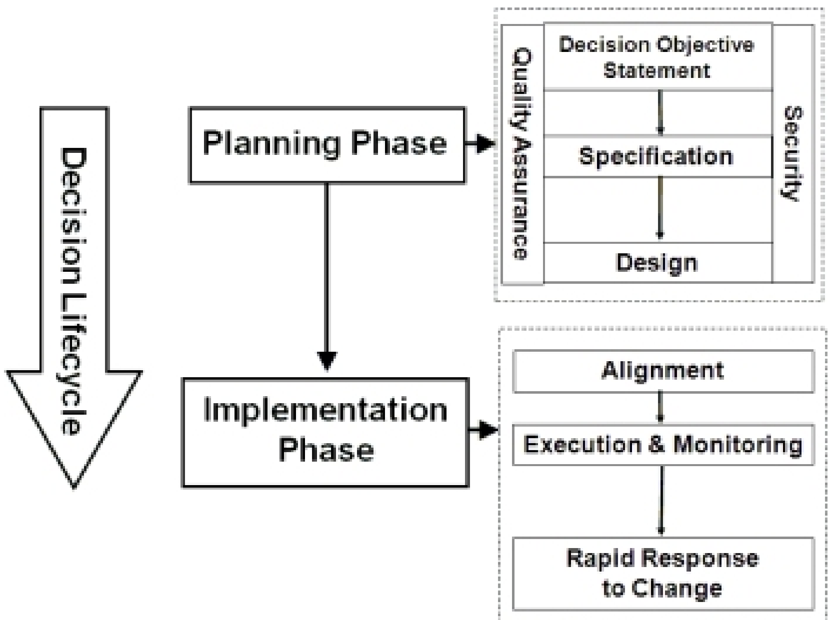

Figure 14: Decision Engineering Framework

There are a number of additional decision engineering stages that complement decision design.

Here, we have illustrated the Decision Objective, Specification, and Design phases of decision engineering. Other steps include the following:

Decision *quality assurance* comprises processes that detect potential errors in a decision. We can evaluate a decision for incorrect assumptions, flawed reasoning, statistical bias, and other errors. Also key to quality is ensuring that uncertainties in the data used to support the decision are clearly identified, and that the effect of these uncertainties, as they propagate through the decision model to the outputs, is clearly understood.

As a decision is *implemented*, *alignment* throughout the organization is critical. This includes ensuring that the decision—including basics like agreement on definitions of key terminology—are consistent throughout the organization.

As decisions change in response to a dynamic environment, more diligent alignment practices become especially important. An organization that requires months to communicate a change in strategic direction to its employees loses agility.

Alignment is assisted by a number of decision engineering best practices, especially the decision design, which helps to communicate the decision and its rationale as the decision planning process evolves, and as the environment changes during the implementation lifetime of the decision.

The core to the implementation phase is decision *execution*, which in all but the simplest of environments must be tightly coupled with *monitoring key assumptions* (as described above) to ensure that they remain inside the bounds within which the decision model yields the desired outcomes. For example, a key assumption in ArwenCom Telecommunications' model might be the "Cannibalization factor" shown in Figure 12. This number might be difficult to estimate during initial modeling, but its value may become apparent as the product is rolled out. Note that dashboards and revenue assurance techniques fit here: instead of monitoring arbitrarily, we do so in the context of planning expectations. Finally, if during monitoring, new information arises indicating the original decision should be adjusted, then decision planning activities are restarted.

Decision Engineering integrates a number of best practices into a single methodology

The above phases provide a structure for organizing decision engineering best practices. Although many existing practices within companies today fit into various parts of this framework, until now they have not been integrated into a single process, nor have they been recognized as constituting an engineering discipline. By unifying them into this framework, we not only have a rigorous approach to planning decisions, maximizing both the quantitative data and human expertise at hand, but a means for utilizing the products of the planning phase to help guide execution of the decision, most particularly, to deal with change over time. Furthermore, by recognizing the methodology specifically as an engineering discipline, we find many opportunities to validate, improve and extend it by comparing it with other engineering disciplines, many of which have aspects of their histories that closely parallel decision engineering.

Perhaps most importantly, decision engineering creates, for the first time, a standardized conceptual framework, with an associated set of processes, to improve the success rates of complex decisions and their implementation. The framework enables decision makers to overcome the complexity of the situations they face, within the timelines available to them, by using analytic techniques that have a long, proven history in other disciplines. Key to the approach is the use of structured visual representations of various aspects of the decision lifecycle, which can be easily shared and unambiguously understood among all decision stakeholders.

Furthermore, the documentation produced during decision planning provides a rich common base for future decisions. Unlike traditional "decision support" materials, which are normally ad-hoc and rarely find on-going use after the decision is made, decision engineering designs have a well-defined structure, and become an integral part of the implementation process. Because they remain as "living" documents and are often re-used, decision engineering processes naturally encode "lessons from history" into an organization's decision making culture.

Spreadsheet Limitations

Spreadsheets, although widely used for decision making, are not a good fit.

It is worth taking a moment here to point out that, in many organizations, a similar activity to the modeling task described in this article is presently done using spreadsheets. Despite the existence of many sophisticated quantitative modeling and decision support tools, the use of such high-end tools is not nearly as embedded in corporate culture as is that of spreadsheets. The reason: these tools are often hard to use, unintuitive, and most importantly, do not speak the language of the decision maker: rather they ask the decision maker to learn a new way to process information with which he is not familiar. Furthermore, the tools are not designed to be modified by the decision maker; rather, this is an (often major) IT project.

Spreadsheets, though in widespread use, are not designed for modeling, and have many drawbacks when used inappropriately in this situation:

Most importantly, in a spreadsheet the **structure of the model is almost completely invisible.** Model relationships are represented typically as textually encoded formulae embedded in cells that link one model element to another; given a model in a spreadsheet, it is usually very difficult to parse the structure of the model unless you are already intimately familiar with it.

To illustrate, consider the following formula from a government spreadsheet that allows consumers to calculate their household carbon footprint:

=I16-IF(I12=1,(F12*52)/(F16+F77)*19.4*(100/95),(F12)/(F16+F77)*19.4*(100/95))

This is extremely difficult to understand, to manipulate meaningfully, to check for errors, and to share with others.

Change management is almost non-existent. It is difficult to understand how changing a relationship in the model in one location propagates throughout the model.

Given the opaqueness of the structure, the model is very difficult to **QA.**

Collaboration and security features are very weak. Just as a set of dedicated tools does not exist for the quantitative modeler, there is also no universally agreed-upon methodology or language for building models or exchanging information between different organizations that may be modeling different aspects of a single complicated system.

Conclusion

Decision engineering supports a decision making team as it works with the complexities of decision making, instead of being overwhelmed by irrelevant data.

This article has shown how a decision model can be constructed to aid a team in choosing whether to launch a new product offering. Although the decision making process is complex using this approach, it focuses effort where it is needed: on the structure of the model instead of on irrelevant data. This approach is also much more tractable than the use of a spreadsheet for the same choice. For many decisions, the cognitive requirements of analyzing the complexities shown here are too great for a spreadsheet. Instead, by using an approach that:

1) Is visual,
2) Organizes information around the *decision* to be made instead of the *data* that drives the decision, much of which may be irrelevant,
3) Allows the decision to be constructed in manageable pieces,
4) Is iterative, and
5) Allows a team to work together to help each other through the complexity,

We see from the example worked here that complex decisions are indeed manageable, and that we do not need to throw up our hands in frustration—as did AveryTel's CEO. Instead, by following a principled methodology, the complexities of decisions with huge consequences—flying the space shuttle or building a skyscraper—can be overcome by a systematic framework.

This approach produces a number of very important benefits over today's typical organizational practices, from the use of spreadsheets to the employment of sophisticated quantitative analysis staff and tools:

1) It is much easier to check a decision model for errors than a spreadsheet

2) The decision model is *owned by the decision making team*, rather than being delegated to a separate simulation, quantitative modeling, or other group (though quantitative modelers are often involved in the Decision Engineering process, their role is around their expertise of data management and statistical (including predictive and simulation) modeling—not substituting for decision maker's judgment.

3) This means that the role of human expert judgment is not separated from the role of quantitative information. These two elements are used *together* within the same formalism.

Decision engineering creates a "navigational" infrastructure for organizations seeking to survive and thrive in complex and dynamic environments.

Although the problem shown here is extensive and time consuming to analyze, for many decisions this effort is well justified. Even so, compared to the time taken using traditional techniques, following the methodology described in this paper when coupled with the appropriate software tools, reduces the decision analysis timeframe to several days, rather than many weeks or even months. It also creates an agile decision making infrastructure that can be adapted as needed. In this sense, Decision Engineering serves as a "navigational infrastructure" for organizations seeking to survive in complex, rapidly changing environments.

---

[1] See Pratt and Zangari: *High Performance Decision Making: A Global Study,* 2008, available by request to inquiries@quantellia.com.

[2] See, for example, the Telco 2.0 web site, at http://www.telco2.net.

[3] See Kerbel, Josh: *Lost for Words: The Intelligence Community's Struggle to Find its Voice,* 2008, at http://www.carlisle.army.mil/USAWC/parameters/08summer/kerbel.htm and Medina, Carmen A., *What To Do When Traditional Models Fail.* Studies in Intelligence, 2004, at https://www.cia.gov/library/center-for-the-study-of-intelligence/kent-csi/pdf/v46i3a03p.pdf.

[4] See, for example, *Laboratory medicine practice guidelines: evidence-based practice for point-of-care testing*. Washington (DC): National Academy of Clinical Biochemistry (NACB); 2006.

[5] Pratt and Zangari, *The New Science of Decision Engineering,* in *Telecom Asia*, September 2008, pp. 20–22. Downloadable from http://www.quantellia.com/Data/TelecomAsia-DecisionEngineering0908.pdf

[6] Pratt and Zangari, Managing Uncertainty, Imprecision and Unreliable Data in Decision Engineering, Decision Engineering Whitepaper, Quantellia [Forthcoming, 2008].

[7] See, e.g. Gigerenzer, Gerd: Gut Feelings: The Intelligence of the Unconscious. 2007.

[8] See, e.g. Feith, Douglas J.: War and Decision: Inside the Pentagon at the Dawn of the War on Terrorism. 2008. Rumsfeld's focus on Key Assumptions as part of the decision making process around 9/11 is reviewed here.