Estimating Benefits of Knowledge Management Initiatives: Concepts, Methodology, and Tools

Joseph M. Firestone, Ph.D.

Knowledge Management Benefits and Corporate Goals

How are various claimed Knowledge Management (KM) benefits related to corporate goals, business processes, and to IT applications? Most discussions of KM benefits and, for that matter, of benefits of other alternatives to KM, are not tightly coupled to corporate goals and business processes [1]. In the KM literature the discussion of benefits thus far has not approached a systematic analysis of corporate goals, objectives and benefits in the context of KM alternatives.

Instead, in most analyses there is an ad hoc listing of envisioned *outcomes* or *effects* of the introduction of KM initiatives and an assertion that these outcomes are unequivocal benefits. The approach is basically intuitive rather than analytical and comprehensive. It doesn't clarify the relationship of the claimed or envisioned outcomes to corporate goals or business processes. And it often doesn't distinguish the outcomes in terms of the degree of benefit they provide.

This paper presents concepts, methodology and tools for producing improved KM benefit estimates. My objective is to provide a framework for thinking about more comprehensive estimation of KM benefits -- estimation that is tightly coupled to corporate goals, and that distinguishes benefits according to their relative importance. I will not propose a specific methodology for estimation in all situations, because, as we will see, no single methodology is appropriate for every corporate situation. Comprehensive benefit estimation is not practical in many situations. While, in others, varying degrees of comprehensiveness will be appropriate.

Instead of a single methodology, I will define an *abstract pattern* of Comprehensive Benefit Estimation (CBE) that would, if implemented, achieve the goal of tight coupling of benefits, goals, and KM initiatives and competing alternatives. Then I will point out how in different concrete situations one may tailor the pattern to achieve a feasible estimation procedure.

A Framework for KM Benefit Estimates

To improve KM benefit estimates both a broader conceptual perspective and more substantial methods and tools are needed than those provided by ad hoc analytical approaches. The road from KM initiatives or programs to benefits leads through business processes and corporate goals. So here is an introductory conceptual framework that can lead us down this road. The first part of the framework relates business processes, corporate goals, and KM initiatives. The second focuses on the relationship between corporate goals and benefits. Once the framework is developed, I will discuss issues related to applying it to estimating KM benefits tightly coupled to corporate goals.

Corporate Goals, Business Processes, and KM Initiatives

Corporate Goals are one category of global property of corporations. Corporate goal-strivings are pre-dispositions to perform actions calculated to create or maintain certain intrinsically valued states of the world, either internal or external to a corporation. Corporate goals are no more than these valued states -- the targets of goal-strivings. I distinguish between corporate goals and corporate objectives by defining objectives as states that are valued instrumentally for the contribution they make toward achieving corporate goals. So there is, in this conception, a cause and effect relation between goals and objectives. Objectives cause an agent to move closer to its goal. Goals may or may not reinforce objectives.

Sidebar One: Analytical, Structural, And Global Properties of Corporations

For every multi-person corporate organization, we can distinguish analytical properties, structural properties, and global properties. [2] Analytical properties are derived by aggregating (summing, averaging, or performing other elementary mathematical operations on) data describing the members of the corporation. Structural properties are derived by performing operations on data describing relations of each member of the business to some or all of the others. Lastly, global properties are based on information about the corporation that is not derived from information about its members. Instead, such properties are produced by the intra-corporate interactions comprising the system they characterize. And, in that sense, may be said to "emerge" from these interactions. [3].

This distinction between goals and objectives is conceptually precise, but actual states of the world may be both goals and objectives. This is true because they can be simultaneously valued in themselves, and for their instrumental value. Corporate goals can be highly abstract, or very concrete. They can also be general in their geographic or temporal focus, or very specific. Of course, highly abstract goals also tend to be very general in scope, while highly concrete goals tend to be very specific. The same variations of abstractness and concreteness and generality and specificity apply to corporate objectives.

Both goals and objectives are often expressed in generalized and vague form in corporate discussions of them. “Our goal is to be the most competitive corporation in our industry.” “Our goal is to be an ethical and socially responsible member of the community.” “Our goal is bring the vision of the integrated desktop to all consumers." These are three examples of vague statements of goals one might find in marketing literature. But, there are also precise ways to express corporate goals. Since goals are states of the world, we can also look at them as sets of ordered attribute values describing the corporation or its environment. Imagine a row in a database table, or a row vector in an algebraic matrix, recording a set of values for a corporate entity. ***This row might define the actual state of the corporation at a particular time.*** Now imagine that this row was made up not of actual values, but of desired values intrinsically valued by a corporation. The row now defines a multi-attribute goal-state of the corporation at the particular time.

The conceptual "distance" between the goal-state and the actual state is *the pre-decision descriptive instrumental behavior gap*. It is the gap that must be closed for the corporation to get to its goal. Figure One illustrates the ideas of the multi-attribute goal and actual states of a corporation through a geometrical interpretation. The geometric space defined by the component attributes of the goal and actual states I will call Corporate Reality Space. The goal and actual states are represented by line vectors drawn from the origin to the points in corporate reality space defined by the attribute values of the components of the vectors. The pre-decision, descriptive, instrumental behavior gap is represented by the distance vector: "a."

A benefit is provided to a corporation when an instrumental action has the effect of moving it closer to its goal-state on one or more of the component attributes of the goal-state. A cost, in the general sense of the term, is levied to the same corporation, if the effect of the action is to move it away from its goal-state on one or more of the component attributes. A net benefit results when the sum of all benefits is greater than the sum of all costs resulting from the action. A net cost results when the sum of all costs is greater than the sum of all benefits. In the geometric interpretation, a net benefit reduces the distance between the actual state and the goal-state. A net cost increases this distance.

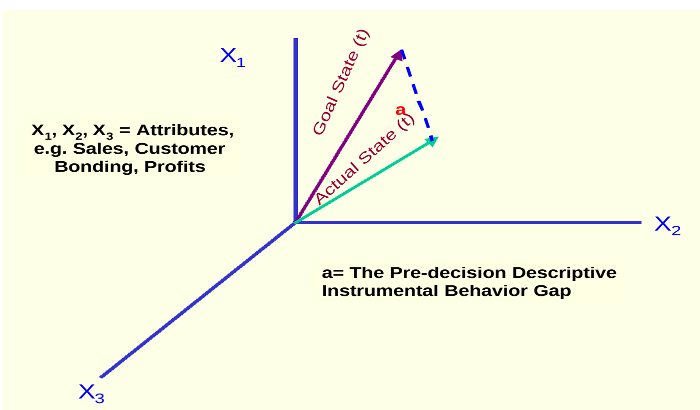

Figure One -- Corporate Reality Space

These statements raise the issue of measurement of the amount of benefit and cost resulting from a decision. While it is generally true that a reduction in the distance "a" can be called a net benefit, the amount of distance reduction is *not* the amount of net benefit. Nor is the amount of increase in "a" the amount of net cost increase. The conceptual distance between the *descriptive* goal and actual states does not, alone, provide enough information to measure *amount* of benefit and cost; *because corporate reality space and the component attributes comprising it are purely descriptive and not evaluative in character.*

To say that there is a net *benefit* when we close the descriptive gap between the actual and goal-states in corporate reality space is to go beyond the purely descriptive character of reality space and to place a *value interpretation* on such movement. But this value interpretation is still less than explicit and somewhat ad hoc, because it assumes a correspondence between reality and benefit without clarifying exactly what this correspondence is. To make the correspondence explicit, we need to work with both a descriptive (corporate reality) representation of goal and actual states and with a valuational (benefit/cost) representation of these. And we need to define a value interpretation mapping corporate reality space to corporate valuational space. I will return to this subject in the section on corporate goals and KM benefits, below.

Corporations try to achieve their goals and to produce benefits by performing business processes. Business process *activities* may be viewed as sequentially linked and as governed by validated rule sets of agents, i.e. their knowledge. [4] [5] [6] [7] A linked sequence of activities performed by one or more agents sharing at least one corporate objective or goal, is a *Task*. A linked sequence of tasks governed by validated rule sets of the agents performing them, and producing results of measurable value to these agents is a *Task Pattern*. A cluster of task patterns, not necessarily performed sequentially, often performed

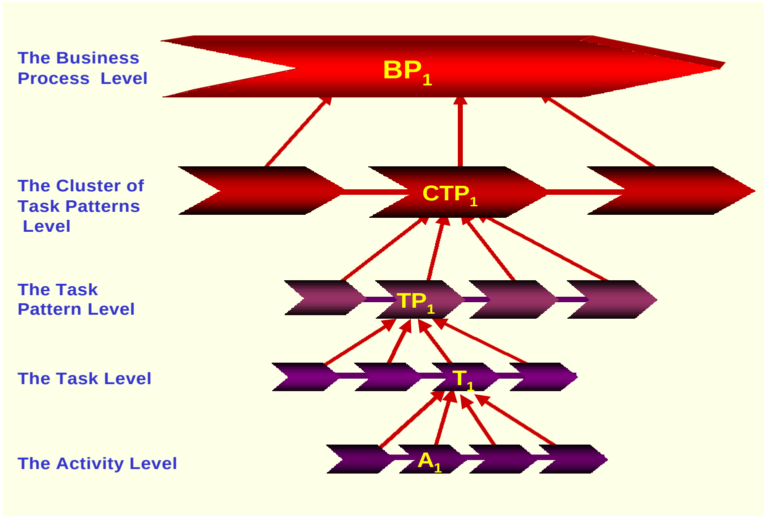

Figure Two -- The Activity to Business Process Hierarchy

iteratively, and incrementally, is a *Task Cluster*. Finally, a hierarchical network of interrelated, purposive, activities of intelligent agents that transforms inputs into valued outcomes, a cluster of task clusters, is a *Business Process*. This activity to business process hierarchy is illustrated in Figure Two.

Business processes in corporations may be evaluated in terms of their efficiency, quality, effectiveness, and net benefit or cost. *Efficiency* refers to the cycle time of the business process compared to some norm. *Quality* refers to how well the activities and tasks constituting a business process are performed relative to some set of quality standards. *Effectiveness* refers to whether or not the business process moves the corporation toward or away from its goals and by how much. *Net benefit and cost* refer to how much a business process is benefiting or costing a corporation.

KM, like other business processes, helps or harms corporations in attaining goals and producing benefits. In order to measure its impact, it is necessary to view it as one of a corporation's business processes, making an impact on other business processes, and, through them, on movement toward or away from corporate goals and/or objectives. In attempting to measure, analyze, or forecast its likely benefits, we need to trace the impact or forecasted impact of the introduction and operation of KM initiatives on knowledge processes. We then need to trace this impact through knowledge outcomes and other business processes, to its further impact on corporate goals and benefits (see Figure Three). Assessments of this kind are not easy or straightforward. But they are necessary if a claim about the likely benefits of a KM project is to amount to more than nonsense or hyperbole.

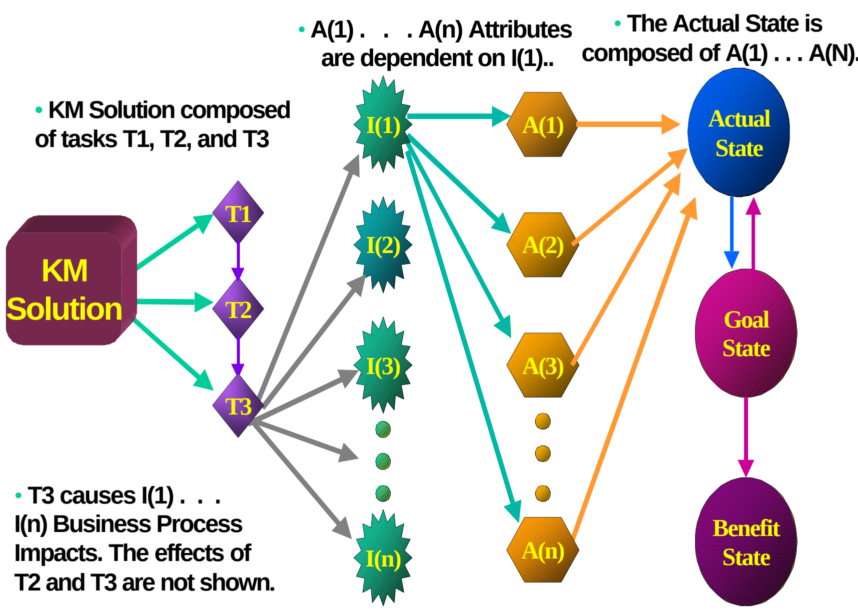

Figure Three – The Path from KM Introduction to Benefit

KM Benefits and Corporate Goals

I pointed out earlier that to relate corporate goals to corporate benefits, we needed both descriptive and valuational representations of actual and goal-states and of the gap between them, and that we also needed a mapping between the two representations. Such a mapping is called a value interpretation. It is a rule (for example an if…then statement), or set of rules that establishes a correspondence between the components of reality space and the components of the valuational space that is the target of the mapping.

From a geometric point of view a value interpretation of corporate reality space is defined by a set of correspondence rules mapping the dimensions (coordinate axes) of reality space onto the dimensions of valuational space. If the valuational space is one whose coordinate axes or attribute components are measured on an absolute benefit measurement scale, then we can call this valuational space *corporate benefit space.*

Both the actual and goal-states will have corresponding vectors in corporate benefit space. Let's call these the *actual benefit vector* and the *goal benefit vector*. The distance between the actual benefit vector of a corporation and the origin of corporate benefit space is the total net benefit enjoyed by a corporation at a point in time. The distance between the goal benefit vector of a corporation and the origin is the total net benefit desired by the corporation. The distance between its actual and goal benefit vectors is the *instrumental behavior benefit gap*. It is this gap, even more than the descriptive instrumental behavior gap that corporations seek to close.

This framework expresses the relationship between corporate goals and benefits clearly. Corporate goals are expressed by the multi-attribute, descriptive, goal-state vector of corporate reality space. Corporate benefits are expressed by the muti-attribute, valuational goal benefit vector of corporate benefit space. The relationship between the two is precisely defined by the set of correspondence rules defining the mapping between the two spaces. Figure Four illustrates the relationship.

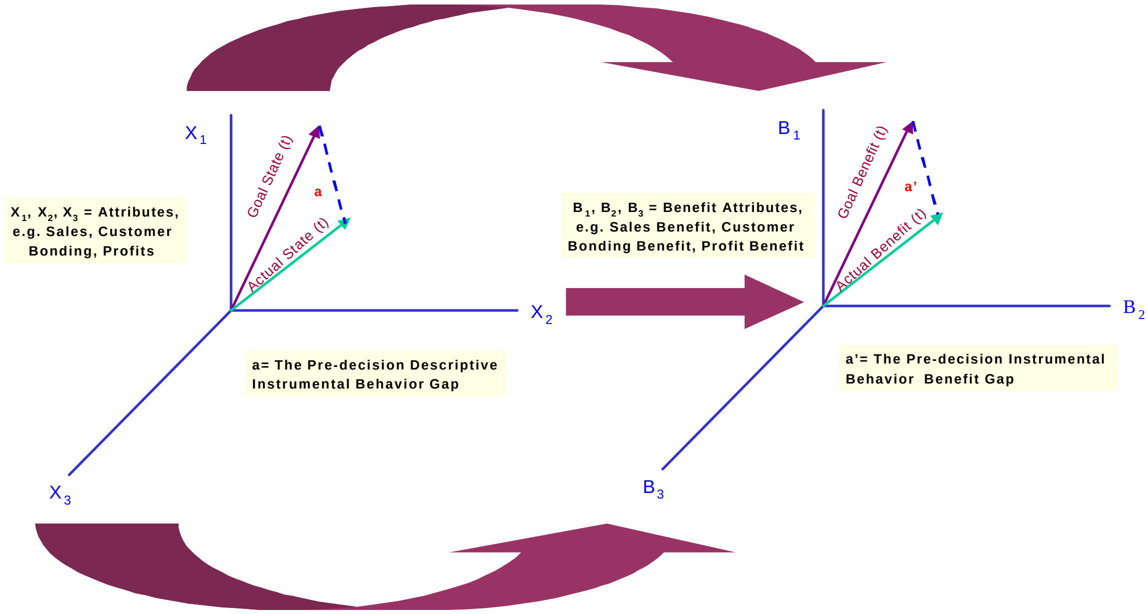

Figure Four -- The Relationship between Corporate Goals and Corporate Benefits

Methodology and Tools for Estimating the Benefits and Costs of KM

If we look at KM benefit assessments from the viewpoint of the conceptual framework, it is clear that a thoroughgoing KM benefit assessment would:

- explicitly postulate and measure goals, objectives and progress toward them
- gauge the impact of KM introduction on business processes and their success in attaining goals and objectives, and finally,
- interpret these descriptive analyses of KM impact or projected impact on goals in terms of corporate benefit so that descriptions of impact are not confused with measurements of actual benefit.

Below is a list of the steps involved in each of these phases of comprehensive KM benefit estimation, along with some comments describing a little of the work involved, and tools that might be used to accomplish it. The list of steps is an outline of a methodology. By comprehensive I mean it takes into account both benefits and costs, and also provides for measuring those benefits and costs that cannot be expressed well in monetary terms.

In this technique of benefit estimation, attributes expressing value in monetary terms are viewed as descriptive attributes. They only become benefit attributes during the mapping from reality to benefit space and after the transformation that mapping entails. So measures in benefit space transcend monetary value and incorporate it into the overall framework. In particular, monetary costs and benefits are measured as a by-product of applying the framework and as a step along the road to more comprehensive measurement.

Measuring Actual and Goal-States

Step One: Perform Measurement Modeling [8][9]

- Conceptualize and select
  - attributes to *describe goals and objectives* in reality space
  - attributes hypothesized to *cause* changes in actual states moving them toward or away from goal-states
  - attributes describing possible unintended *side effects* of actions activating causal attributes
  - other *outcome* attributes important for description

Many of the selected attributes will be abstractions. These are not defined by data attributes, but are assumed to be computable from them. They are measurable attributes. Other attributes are directly defined by data, and are measured attributes. So we have a mix of measured and measurable attributes in each of the four categories.

- Organize attributes into measurement clusters

That is: group the abstract attribute (e.g., customer acquisition, customer retention, customer profitability, revenue growth) that is the target or focus of measurement, with the set of already measured attributes that will provide values to be used to compute, measure, or derive the values of the abstract target attribute. The outcome of this task is a categorization of measured attributes by the measurable attributes that are the primary focus of the measurement modeling effort.

- Construct measurement models

These are models made of rules expressing measured attribute values as antecedent conditions and target (measurable) attribute values as consequents, with no temporal priority specified between the antecedent and consequent values. They are not limited to goals and objectives, but include causes, side effects and other outcomes as well. The rules are called semantic rules or rules of correspondence. They frequently create multi-attribute composites as measures of the target attributes. Such composites can often be complex and demanding to construct. Measurement models are essential to modeling, and can always be distinguished as a logical component in any systems model. You can't formulate a testable model of an aspect of the world without using a measurement model. The only question is whether the measurement model is explicit or implicit.

There are at least four types of ("crisp logic-based") measurement rules that provide the foundation for a composite in measurement models. [9] In addition, there are fuzzy logic-based rules that map crisp values to linguistic variable values through fuzzy membership functions. [8] Among other things, such rules establish priorities among the attributes entering the composite.

The activity of prioritizing attributes for their relative importance to a criterion variable is frequently part of measurement modeling. A measurement model is different from a causal model, in that the latter requires temporal asymmetry between antecedent and consequent [10], while measurement models imply that values of the measured and measurable attributes are being viewed in cross-section. I've provided detailed accounts of measurement modeling in other places, and won't review them here. [8] [9]

- Use ratio scaling methods where possible

Ratio scaling methods should be used in constructing measurement models of abstractions because they allow easy mathematical manipulation and composition of measured attributes into target attributes. That is, they're easy to work with when you want to create an overall measure from a set of component attributes.

Ratio scaling should be used in doing priority assessments among attributes, because the resulting weights are defined on the same scale and facilitate combining the measured attributes into the target attributes. It should also be used along with direct judgmental assessments of quantitative properties of agents and corporations where attributes are not measured and you want to convert them to measured attributes. In that case ratio scaling provides a further basis for combining measured attributes into target attributes.

Ratio scaling techniques are now well-known and easy to implement. Saaty's work on the subject is particularly accessible, and his development of the Analytical Hierarchy Process (AHP) over a period of more than thirty years has featured an emphasis on practical ratio scaling methods and their application to a wide variety of subjects. [11] I have also recently treated techniques for ratio scaling in the context of knowledge discovery in databases. [8]

Step One can be greatly facilitated by the proper tools. While a wide variety of tools including a spreadsheet such as Excel with fuzzy logic add-ins, a mathematics package such as MATHEMATICA or MATLAB (with Fuzzy logic add-ons) can accomplish everything you need to do in this step, the best combination of ease of use and power will be found with an omnibus statistical/modeling package such as SAS, SPSS, or Statsoft (my personal favorite because of its great spreadsheet module, graphics and general combination of friendliness and power), supplemented by Inform Software's Fuzzytech. With these two tools you can accomplish all of the measurement modeling and ratio scaling you need to do, and when you're done you can communicate results to external packages and back-end databases.

Step Two: Gather Data or perform direct assessments to measure attributes that cannot be derived from measured attributes or that have no data

Once the measurement model is constructed, you can't apply it without having values for its measured variables. You get these by gathering data from documents, surveys, or direct observation. Data from these sources is preferable to data gathered from direct assessments of properties of corporations and agents by "expert raters" because it is thought to have greater reliability and validity. But data gathered from direct judgmental assessments produced by panels of expert raters is certainly better than missing data. And there is a good deal of evidence in both Saaty's work [11] and some of my own [12] suggesting that reliability and validity levels in models using direct assessment data are comparable to those achieved by models using document or survey-based data. And that they exceed the levels found in models based on opinion and attitude surveys.

Step Two uses the same tools as Step One for direct assessments. For more conventional data gathering and data staging you may need data warehousing ETL/data cleansing tools such as Informatica, Sagent, Informix's Datastage, or Evolutionary Technologies" ETI-Extract, and, of course, a commercial relational database.

Step Three: Determine Actual States by using measurement models to compute attribute values

Once values are given to the measured attributes, the measurement model is used to compute the values of the target measurable attribute(s) to arrive at a description of the actual state.

Step Four: Determine Goal-States by specifying goal attribute values

The goal-state could be specified without first determining the actual state. But it is easier to do a complete job of specifying the goal-state once the actual state is measured and available for examination. Then one can begin by using the actual values of the attribute components of the goal-state as a baseline for estimating the goal-state values of the same attributes. A variety of methods can be used to perform these estimates, including pair comparison rating methods. The trick is to estimate goal-state values at different points in the future, so there is enough data to measure the logical consistency of the judgmental estimates. Once these values are derived from the estimation procedures a consistency check can also be made on the fit between the computed future values of the abstract attributes and the judgmental estimates of those abstract values. The judgmental forecasts of measured attribute values may then be adjusted until they are consistent with the forecast values of target attributes.

Step Five: Compute the pre-decision instrumental behavior gap

Subtract the goal-state vector from the actual state vector to get the distance or gap vector. Compute the length of this vector, which, in the most commonly used mathematical interpretation, is the euclidean distance between the goal-state vector and the actual state vectors.

Steps 3, 4, and 5 require no additional steps.

Modeling the Impact of KM Solutions

Step One: Select abstract attributes that are the focus of measurement models as target attributes for Impact Modeling

Classify these attributes into exogenous attributes, mutually endogenous attributes, and endogenous attributes. The *exogenous attributes* are causes of other attributes, but are "not caused " by any other attributes included in the impact model. *Mutually endogenous attributes* have effects on other attributes in the model and are affected by these same attributes. *Endogenous attributes* are affected by other attributes and only affect other attributes without being affected by them. No additional tools are needed for this step.

Step Two: Specify Impact Model

Specify hypotheses expressing the values of:

(1) mutually endogenous attributes as a function of other mutually endogenous attributes and exogenous attributes, or as a function of exogenous attributes alone;

(2) mutually endogenous attributes as a function of other mutually endogenous attributes; and

(3) endogenous attributes as a function of mutually endogenous attributes.

In these hypotheses all determining attribute values must temporally precede all determined attribute values.

The result of this step has been variously called a cognitive map, a conceptual graph, a causal model, an impact model, a semantic network and many other terms. It is composed of nodes and connecting rules. The rules can be expressed in terms of "crisp" logical rules, or in terms of fuzzy rules. If the latter is true, it has been called a "Fuzzy Cognitive Map." [13]

A variety of modeling tools can be used for this task depending on one's impact modeling orientation. The statistical packages mentioned earlier support *linear structural equation modeling*. In the *system dynamics* area good choices are Ventana's Vensim, and High Performance Systems, ithink and Stella products. Another leading product is Powersim. For *Fuzzy Cognitive Mapping* you can use Fuzzytech. If you prefer a *complex adaptive system* (cas) approach, then Santa Fe Institute's agent-based Swarm simulation, is the indicated choice.

Step Three: Expand the impact model by adding hypotheses comparing the effects of KM and other software alternatives on mutually endogenous and endogenous variables

It is better if this is done by expressing the software alternatives in terms of component attribute values that describe them and allow (1) direct comparisons of features among alternatives, and (2) formulation of hypotheses relating software features to business process attributes. However, you can also add hypotheses specifying the relative magnitude of the impact of KM software versus alternative software options on each mutually endogenous or endogenous attribute in the model. Ratio scaling techniques can also be used here to measure these relative magnitudes and to check on the consistency of the judgments. No additional software is needed for this step.

Step Four: Implement empirical tests and simulations of competing impact models and evaluate software alternatives

The tests will provide forecasts and analyses of the impact of KM vs. other types of solutions in moving the corporation toward its goal-state. If you use a Fuzzy Systems approach to impact modeling you'll need some neural network estimation software for testing and validation. The statistical packages mentioned above also provide such software, as does Fuzzytech, which specifically supports neuro-fuzzy estimation. But if you're willing to go beyond these packages to vendors more specialized in neural networking, Ward Systems Group and NeuroDimension offer excellent and versatile software.

Mapping from Reality to Benefit Space

Step One: Define rules of correspondence between attributes of reality space and attributes of benefit space

There are a number of things to keep in mind when doing this mapping. First, only some of the target attributes of reality space need be directly represented by attributes of benefit space. The determining factor is whether an attribute is intrinsically valued as a benefit. An important implication is that benefit space may be of much lower dimensionality than reality space. Second, an attribute of benefit space may be the result of a composite mapping from multiple attributes of reality space. This is another source of possible lower dimensionality in benefit space. Third, even if an attribute in reality space is represented in benefit space, the mapping is unlikely to be a simple correspondence in values. Mappings can be similarity transformations, linear transformations, non-linear transformations of various kinds, and fuzzy membership functions of diverse form. [14] A mapping from an attribute in reality space to a corresponding attribute in benefit space is called a principle of correlation. [15] [9] Such a rule should be validated through consistency testing and graphical means [9].

Step Two: Establish benefit priority weights among attributes of benefit space

This is done using the same type of ratio scaling techniques used in measuring actual and goal-states.

Step Three: Compute the instrumental behavior benefit gap

The same method can be used as in computing the descriptive instrumental behavior gap, but keep two differences in mind. First, the attributes used are mapped transformations of the descriptive attributes called benefit attributes. And second, in computing the euclidean distance, priority weights determined in step two are used to weight the attributes of benefit space to arrive at an overall measure of benefit. No additional software is needed for these three steps.

Implementing Estimation

The estimation methodology I described has the advantages of being comprehensive, and of tying the analysis of benefits to corporate goals, but the disadvantage of being expensive in effort and money. It is much more likely to be used to evaluate a KM initiative after the fact than it is to be used to forecast likely impact during the planning stage. Not least because, if begun from scratch, it will take months to implement, an unacceptable time period for a KM planning study. To make it useful then, abbreviated versions of it are needed that will represent an improvement over ad hoc benefit analysis, but that can still be accomplished in a few weeks of effort. The nature and extent of abbreviation will depend on the corporate environment encountered. Here are three cases. Together they define the limits for abbreviating the methodology. Real world situations will fit some synthesis of the cases.

Case One: No prior work on development of an Enterprise Performance Management (EPM), balanced scorecard, ERP, or data warehousing system

This situation is hard to imagine in any major corporate environment today. In it, the comprehensive methodology of benefit estimation cannot be applied without going through all of the steps outlined, because little prior work on measurement and impact modeling already exists. In a situation like this one, it won't be possible to accurately estimate the benefits of a KM initiative relative to another KM, or data warehousing, or balanced scorecard system, with any degree of confidence without months of effort.

On the other hand, this is also the situation where a KM initiative, or any other competing alternative, is likely to have its highest ROI; because it introduces a whole system of measurement and performance analysis which was previously not available. So a decision selecting any data warehousing, EPM, ERP, balanced scorecard, or KM initiative can be made with reasonable confidence of substantial payback.

Once this point is recognized, the question becomes not so much whether KM produces enough benefit that it should be funded, but whether a specific KM initiative should be funded in preference to one of the other alternative initiatives that can improve knowledge production, and delivery to end users. If the question is whether a KM initiative will bring greater benefits than other alternatives, rather than the broader one of providing an estimate of KM benefits relative to those provided by other alternatives, then there is an inexpensive method of benefit assessment that can be used to project the impact of various program alternatives relative to one another.

The method is Saaty's Analytical Hierarchy Process (AHP). It has been applied by Fatimeh M. Zahedi to quantitative evaluation of expert systems [16], and can be adapted to the problem of deciding which of a group of KM or other program alternatives will provide the greatest benefit relative to other members of the group.

The nice thing about using the AHP when no prior work is available is that it needs no measured data to work except data generated by the method from judgmental assessments. It takes judgmental assessments about decision options generated at the lowest, most concrete level of a hierarchy, and combines that data with ratio scaled attribute priority data also generated by judgmental assessments. Let's review the method in detail.

The Analytic Hierarchy Process (AHP) and KM benefit estimation

It is easy to develop a set of criteria to use in comparing alternative programs. One of these can be cost. You can rate alternatives according to monetary expense. Find out which is the most expensive, which the least, and which alternatives are in the middle. You can go on to compare alternatives on other criteria of evaluation, presumably criteria relevant to non-monetary costs and to benefits. But when your comparisons are all done, how do you assess the degree of non-monetary cost? The degree of benefit? How do you combine different non-monetary costs to arrive at summary measures? How do you combine monetary and non-monetary costs? How do you combine different benefits to arrive at summary measures? How do you compare costs and benefits to arrive at benefit/cost ratios? More generally, what is the relationship between specific costs and the overall summative concept of cost? What is the analogous relationship on the benefit side?

None of these questions can be answered by an ad hoc comparative approach. If you use one, the process you use to aggregate comparisons on individual criteria into an overall assessment of which KM or another alternative program is best for you and your organization, will of necessity have to be a subjective, implicit process. The justification for your choice will certainly be incomplete, probably flawed, and subject to obvious criticism from those who quickly perceive other subjective criteria you didn't consider.

A better comparative evaluation is produced if you use a rigorous framework that (a) specifies the meaning of benefit and cost in terms that *connect and tightly couple the overall goals of your organization and the characteristics of the KM solution and other alternatives* within the context of a broad benefit/cost concept. And (b) provides a means of *quantitatively comparing goals, characteristics and intermediate criteria* comprising the evaluation scheme. In such a framework, the criteria used to directly assess the alternative systems are themselves assessed by other criteria that are more directly related to cost or benefit. Then these are assessed relative to still other criteria more directly related to cost or benefit, and so on, until one reaches a set of criteria that may themselves be directly evaluated in terms of overall cost or benefit. The last step in this progression produces a simplified mapping from reality to benefit space. The fact that all the assessments in the progression are quantitative in nature means that questions of the sort posed above may all be answered by applying the evaluation framework.

In addition, any criticisms involving formulation of additional evaluation criteria to be applied to the alternatives would have to be related to the benefit/cost framework before their validity could be asserted. If they were so related, further, they would not change the result of an evaluation unless their quantitative significance was great enough to have a major impact on overall scores. In other words, in contrast to subjective evaluation frameworks, a rigorous evaluation framework of the kind offered here produces cumulative results. Even if mistakes of omission are made, the results of a prior evaluation need not be scrapped but only revised, and the overall result is much less likely to be disturbed.

The AHP fits the specification for a rigorous comparative evaluation and assessment framework just described. The AHP has had almost 30 years of development, since its inception in the early 1970's. The primary developer of the

AHP is Professor Thomas L. Saaty of the University of Pittsburgh. Saaty began work on certain aspects of the AHP while he was with the U.S. Arms Control and Disarmament Agency (ACDA). He published the first studies applying the AHP during the early 1970's, when he had moved to the University of Pennsylvania. [17] [18] There, and later at the University of Pittsburgh, Saaty, his colleagues and students have applied the method to a wide range of practical problems, including planning, prioritizing, optimization, benefit/cost analysis, decision making, the study of national influence, terrorism, international conflict, transportation, energy policy, and many other areas. [19] [20] [21] [22] These, represent only a few of a voluminous list of references available [11].

The AHP was developed by Saaty to provide a rational basis for multi-criteria decision making of the sort involved in evaluating and selecting program alternatives. The AHP has three aspects. First, it focuses on decomposition of the decision problem to identify various components of an attribute hierarchy -- goals, objectives, criteria, sub-criteria, elements, actors, or characteristics -- relevant for decision; and grouping and ordering these components into sets or clusters of attributes, comprising the levels of a hierarchy, or into clusters that share levels of a hierarchy.

Second, it focuses on comparative judgments between pairs of attributes within each cluster. These comparative judgments serve as the raw data for computations producing a priority rating of each component attribute of a cluster compared with all other components of that cluster in relation to some criterion attribute specified at the immediate higher level of the hierarchy. The priority ratings produced are defined on a ratio scale. One ratio scale is defined for each criterion attribute. Since the priority ratings are defined on a ratio scale, they can be meaningfully multiplied or divided, thus providing a basis for later benefit/cost computations. The ratings are also tested for logical consistency within the AHP, so the extent of departure from the classic ratio scale model is measured and used to evaluate the validity of AHP evaluations.

Third, the AHP focuses on the synthesis of priorities within the hierarchical framework. This means that "local" priority ratings, those established for a particular component attribute in relation to other components of its cluster and a particular criterion variable, are adjusted or weighted according to priority ratings computed for that component in relation to other criterion variables. It also means they are weighted by the priority ratings computed for the criterion variable itself within its hierarchical level and its cluster.

This adjustment process results in a "global" priority rating being determined for each component of each cluster and level of the hierarchy relative to the focal concept or goal of the hierarchy. If the focal concept is a goal attribute such as customer profitability or monetary cost, and the decision alternatives at the bottom of the hierarchy are KM and other program alternatives, then the analytic hierarchy provides no mapping from reality to benefit space. It only provides a kind of relative impact evaluation of each alternative on the goal attribute. If the focal concept (or concepts in the case of more than one goal attribute) is a benefit attribute, then the priority weights defined at the goal level of the hierarchy (relative to the benefit level) will provide a simple mapping (a linear composite) from reality space to benefit space, as well as a relative assessment of the impact of the program alternatives on goals and benefits.

Thus, in a benefits hierarchy the adjustment process produces a rating or measure of the "global" or overall contribution to “benefit” of each member of the hierarchy. The analogous result also holds for a cost hierarchy. So, if global ratings of program (or any other decision) alternatives at the lowest level of an attribute hierarchy are developed through the synthesis of priorities, the AHP yields benefit and cost ratings that may then be divided to arrive at meaningful benefit/cost ratios for each KM or other program alternative being evaluated. As with local priority ratings, consistency tests for "global" priority ratings are also provided by the AHP. If observed inconsistency is too great, the hierarchy may be revised until consistent ratings are provided by decision makers.

Software for implementing the AHP has been developed by Ernest Forman and Tom Saaty at Expert Choice, Inc. [11] The Expert Choice software comes in a number of versions for individuals and teams and enterprises. In all cases it is friendly and it is much easier to implement the AHP with it then it would be using the computational tools recommended earlier for occasional prioritizations using the AHP method.

Case Two: ERP and/or Data Warehousing Systems Exist

This situation is more favorable for implementing the comprehensive benefit estimation methodology. Much of the work of specifying attributes in the goal/objective, cause, side effects, and outcome categories will have been done.

In addition, data will have been gathered on many of the attributes in the system. Still, comprehensive benefit estimation will remain difficult, because most of the steps in the measurement, impact modeling, and mapping of reality to benefit space categories will remain. The solution is to once again apply the AHP, but to substitute measured attribute values in the hierarchy where they are available, and to use the real data to enrich AHP judgments where necessary.

Case Three: A Balanced Scorecard or Enterprise Performance Management (EPM) System is already available

This is the most favorable situation for implementing comprehensive benefit estimation prior to KM construction. Such systems contain measured attributes, measurable attributes, goals, objectives, causal and side effect attributes, and outcome attributes. Balanced scorecards also include cause and effect hypotheses relating measurable attributes to one another.

To implement comprehensive benefit estimation on a balanced scorecard or EPM foundation, supplement the measurement and causal models already present with additional rules and hypotheses, particularly those relating program alternatives to mutually endogenous and endogenous attributes in the model. In addition, the presence of "missing measurements" will also require generating data from direct judgmental assessments. Nevertheless, most of the data gathering, measurement and causal modeling activities will have already been completed. The main remaining task is the mapping from reality to benefit space. This mapping should follow the pattern for the general method I described earlier.

Summary

This paper presented concepts, methodology and tools for producing improved KM benefit estimates. It provided a corporate reality/corporate benefit space framework for thinking about more comprehensive estimation of KM benefits -- estimation that is tightly coupled to corporate goals, and that distinguishes benefits according to their relative importance. No single methodology is appropriate for every corporate situation. Comprehensive benefit estimation is not practical in many situations. While, in others, varying degrees of comprehensiveness are appropriate.

Instead of a single methodology, the chapter defined an *abstract pattern* of Comprehensive Benefit Estimation (CBE) that would, if implemented, achieve the goal of tight coupling of benefits, goals, and software alternatives. It then showed how the general pattern could be abbreviated and tailored in three different "ideal type" situations to achieve a feasible estimation procedure. Actual situations will mix the characteristics of these ideal types

References

[1] Thomas. H. Davenport, and Lawrence Prusak, Working Knowledge: How Organizations Manage what they Know (Boston, MA: Harvard Business School Press, 1998).

[2] Paul F. Lazarsfeld and Herbert Menzel, "On the Relation Between Individual and Collective Properties," in Amitai Etzioni (ed.), Complex Organizations (New York: Holt, Rinehart and Winston, 1961)

[3] John H. Holland, Emergence (Reading, Mass.: Addison-Wesley, 1998)

[4] Edward W. Swanstrom, "What is Knowledge Management?" available at http://www.km.org/introkm.html.

[5] Joseph M. Firestone, "Enterprise Knowledge Management Modeling and Distributed Knowledge Management Systems," Executive Information Systems, Inc., Wilmington, DE, January 3, 1999, available at http://www.dkms.com/White_Papers.htm.

[6] Joseph M. Firestone," The Artificial Knowledge Manager Standard: A Strawman," Knowledge Management Consortium, International, Gaithersburg, MD, January 25, 1999 available at http://www.dkms.com/White_Papers.htm.

[7] Knowledge Management Consortium, "What is Knowledge Management? A Complex Adaptive Systems Approach," KMCI PowerPoint Presentation, Draft 3.0, February, 1999

[8] Joseph M. Firestone, "Knowledge Management Metrics Development: A Technical Approach," at http://www.dkms.com/White_Papers.htm.

[9] Joseph M. Firestone and Richard W. Chadwick, "A New Procedure for Constructing Measurement Models of Ratio Scale Concepts," International Journal of General Systems, 2 (1975), 35-53.

[10] Joseph M. Firestone, "Remarks on Concept Formation: Theory Building and Theory Testing," Philosophy of Science, 38 (Dec. 1971), 570-604

[11] http://www.expertchoice.com

[12] Joseph M. Firestone and Sidney Brounstein Strategic Evaluation and Planning System (STEPS): The Needs Assessment Capability (NAC) - A Description of Products (Co-authored with Sidney H. Brounstein). Program Evaluation Staff, Farmers Home Administration, USDA, September, 1981, Pp. 127-151.

[13] Bart Kosko, Neural Networks and Fuzzy Systems (Englewood Cliffs, NJ: Prentice-Hall, 1992)

[14] Earl Cox, The Fuzzy Systems Handbook (Cambridge, MA: Academic Press, 1994).

[15] Brian Ellis, Basic Concepts of Measurement (Cambridge, England: Cambridge University Press, 1966), P. 41

[16] Fatemeh M. Zahedi. Intelligent Systems for Business: Expert Systems with Neural Networks (Belmont, CA: Wadsworth, 1993). Chapter 12.

[17] Thomas L. Saaty, "An Eigenvalue Allocation Model for Prioritization and Planning," Energy Management Policy Center, University of Pennsylvania, 1972.

[18] Thomas L. Saaty, "Measuring the Fuzziness of Sets," Journal of Cybernetics, 4, no. 4, (1974), 53-61

[19] T. L. Saaty and Shubo Xu, "Recent Developments in the Analytic Hierarchy Process, " in Thomas L. Saaty, The Analytic Hierarchy Process: Panning, Priority Setting, Resource Allocation, 2nd Edition (Pittsburgh, PA: RWS Publications, 1990)

[20] P. T. Harker (ed.) "Special Issue on the Analytic Hierarchy Process," Socio-economic Planning Sciences, 20, no.6 (1986)

[21] F. Zahedi, "The Analytic Hierarchy Process: A Survey of the Method and Its Applications," Interfaces, 16, no. 4 (1986) 96-108.

[22] Thomas L. Saaty. Decision Making for Leaders (Pittsburgh, PA: RWS Publications, 1990)

Biography

Joseph M. Firestone, Ph.D. is Vice-President and Chief Knowledge Officer of Executive Information Systems (EIS), Inc. Joe has varied experience in consulting, management, information technology, decision support, and social systems analysis. Currently, he focuses on product, methodology, architecture, and solutions development in enterprise Information and knowledge portals. He also performs knowledge and knowledge management audits, training, and facilitative systems planning, requirements capture, analysis, and design. Joe was the first to define and specify the Enterprise Knowledge Portal Concept. He is widely published in the areas of Decision Support (especially Enterprise Information and Knowledge Portals, Data Warehouses/Data Marts, and Data Mining), and Knowledge Management, and has written a full-length industry report entitled "Approaching Enterprise Information Portals."

Joe is a founding member of the Knowledge Management Consortium International (KMCI), Its Secretary, a member of its Executive Committee, of the Governing Council of the KMCI Institute, the Director of the KMCI Research Center, the Director of the Certified Knowledge and Innovation Manager (CKIM) Program and the Editor-in-Chief of the new journal "Knowledge and Innovation: Journal of the KMCI." Joe is a frequent speaker at national conferences on KM and Portals. He is also developer of the web site www.dkms.com, one of the most widely visited web sites in the Portal and KM fields. DKMS.com has now reached a visitation rate of 125,000 visits, and 950,000 accesses annually.