

Reducing Risk by Killing Your Worst Ideas

Joseph M. Firestone, Ph.D. Managing Director, CEO Knowledge Management Consortium International

By permission (copyright) Joe Firestone

A Little Background

In addition to consulting work on portals, succession planning, and KM strategy, and teaching Workshops, I continue to develop my ideas on the scope of the field of Knowledge Management and its relationship to other types of management and to business processes and decision making. In the last five years, or so, I've seen Knowledge Management experience what I call "conceptual drift," a subject I've addressed in a number of other places, but, first, within AOK in a debate a few years ago with Dave Pollard.

Conceptual drift refers to the phenomenon of identifying "KM" with one of the techniques, tools, or areas of focus that have an important or very productive place in some KM projects. It involves, essentially, substituting the means we sometimes use to implement KM interventions or knowledge processing with the end of doing KM itself, so that Communities of Practice, or Portals, or Social Network Analysis, or Content Management, or other areas of focus, are thought of as KM in themselves.

Watching conceptual drift play out, it seemed to me that one of the reasons why it was easy for other fields to "colonize KM" was that KM's conceptual core was weak, in that there was no agreement and common understanding among us on our key terms, approaches, and principles. Discussions among KM practitioners and theorists exchanging views on core concepts didn't seem to be helping much in creating such a conceptual core. In fact, there was a surfeit of ideas and frameworks, but a shortage of agreement on how to evaluate and select among the various frameworks, with the result that little selection occurred.

Meanwhile with each passing year, new frameworks were created, adding to the conceptual variation, and agreement on a conceptual core seems further and further away. So, when do we get to the selection and retention to go with the variation? When does our creativity in developing new ideas, begin to pay dividends in providing us with frameworks that actually begin to crystallize what our discipline is about?

Only, I think, when some of these frameworks begin to produce value for us in specific fields of application that people really care about. Here, as in all fields, "the proof of the pudding is in the eating." And KM's conceptual core will emerge from competition among frameworks in terms of how successful they are in application contexts. People will come to understand, and will take on board, a few frameworks, only when they have some very strong and obviously successful exemplars embodying those frameworks.

This brings us to the question of application areas for KM. A few years ago, when the reasoning I've just related was running around in my head, I began to look for a fairly general application area. The topic of our discussion in this installment of the AOK Star Series, "Reducing Risk by Killing Your Worst Ideas," is about applying KM to such subjects as "risk," "risk management," "risk intelligence," and "risk intelligence metrics." It is the application area I selected, where, I hope, the value of a particular KM/knowledge processing conceptual framework can be demonstrated. The conceptual, framework involved is the one I've been developing for years and for those of you not familiar with it, or familiar with only a part of it. You'll find lots of reading material here:

http://www.kmci.org/kmci_resources_and_links.html

<http://radio.weblogs.com/0135950/2004/03/20.html>

This framework includes the ideas associated with the term "KMCI Knowledge Life Cycle," but it also goes way beyond the KLC and is comprised of a number of inter-related frameworks of which the KLC is only one.

I first became interested in the topic of risk as an area of application for Knowledge Management several years ago, when Simon Lelic interviewed me for an article he was writing for Knowledge Management (before it became Inside Knowledge) called "Managing Knowledge to Manage Risk" (Knowledge Management, Vol. 6, No. 1 (2002), 10-13. Later I benefited from reading Dale Neef's *Managing Corporate Reputation and Risk*, Elsevier, 2003 and also his contribution to the special issue of *The Learning Organization* Mark McElroy and I edited, called "Has Knowledge Management Been Done," which contained an article by Dale reprising some of his fine book.

Finally, in developing my ideas on risk, I've benefited greatly from many talks with my friend David Apgar, author of *Risk Intelligence*, Boston, MA: Harvard Business School Press, 2006. Our talks focused on his book, our ideas about value theory, the nature of risk, and our common appreciation of Popper's writings and Critical Rationalism more generally. Our differences in approach helped me to clarify my own views about, and to understand much more clearly, how my approach differs from one much closer to Bayesianism and to Shannon's ideas on information. David's work also led me to formulate an alternative model to his Apgar Test for competitive risk intelligence.

"Reducing the risk of error by killing your worst ideas" is about risk intelligence and risk intelligence management and the connection of these to Knowledge Management. It's also about risk intelligence metrics and approaches to interventions for helping us to reduce the risk of error in our decision models. You can think of these as KM

interventions for increasing risk intelligence and reducing the risk of error.

Basics on Risk and Closely Related Matters

I'll begin my account about how to reduce risk, by pointing to two different categories of it.

- First, there is the risk arising out of a decision; and
- Second, there is the risk of error in one's decision model.

The risk accompanying a decision arises because most major business decisions don't immediately cause or bring about their intended outcomes. Instead, decisions only create a propensity, a probabilistic disposition, for many of their most important outcomes to occur.

Sidebar 1: Dispositions, Predispositions, Propensities, and Probabilities

From time-to-time I'll use words like 'disposition,' 'predisposition,' and 'propensity,' in connection with 'probability' to refer to the tendency of objects or situations to exhibit certain properties or characteristics under specified conditions. Let's look for a moment at some of the distinctions between these terms and also how they relate to one another.

A disposition is a property of an object that shows itself only under certain specified conditions. A favorite example is a magnetic object. When you see one under 'normal' conditions, you can't tell it's magnetic. But place iron filings in its vicinity, and you can tell, because it will attract them.

Sometimes people will talk about 'predispositions' rather than 'dispositions' to emphasize that certain properties will predispose individuals or groups to behave in particular ways when they're exposed to certain conditions or situations. Mostly, such predispositions don't determine that an individual will behave in a particular way, but only impart a tendency to behave in that way given the conditions involved. A very visible and significant example of a predisposition, both to believe certain things, and to behave in certain ways, is racial prejudice. Another is religious fundamentalism. Yet another, more significant in the context of risk, is 'risk aversion.' The study of psychology is, in large part, the study of predispositions. And when we talk about our mental beliefs, we are most often talking about predispositions, rather than beliefs we are actually experiencing at any moment.

Taking a perspective broader than individual or group behavior, whenever there is a tendency for something to occur based on certain specified generating or background conditions, we can say that there is a 'propensity' for that something to occur. It turns out that this notion of 'propensity,' as a tendency for something to occur, given a set of generating conditions, is one of the coherent ways in which the abstract mathematical idea of 'probability' can be interpreted. We can view 'probability' as a weighting of 'possibility' according to the real world tendency of a possibility to occur, given a set of generating conditions.

This, of course, is not the only, or even the most popular way of interpreting the mathematical idea of probability, but it is, for two important reasons, the one I will use here when I refer to real probability rather than the probability of error in one's decision model. The first reason is that it is an objectivist interpretation of probability implying that some probabilities may not be due simply to our lack of information, but to the way the world is. And secondly, I believe that that interpretation, among the various reasonable objectivist alternatives, best fits the important idea that when we make decisions and take actions, the ensuing outcomes are most often not determined by our actions, but are only constrained or influenced by them. In other words, our decisions and actions affect the propensities or probabilities that certain outcomes will occur.

Consider a decision and its outcomes (See Figure 1). Some of these have positive value for us, while others have negative value. The decision has a probability relation to each of its possible outcomes. The probability is the strength or propensity of the decision to bring about the outcome. The combination of the probability and the value of a negative outcome is the risk associated with that outcome, given the decision. So:

== there is an individual risk associated with each negative outcome and its probability, and

Figure 1: A Flattened Decision Risk Model

== the total risk of making the decision, provided the outcomes are independent of one another, is the combination of all the individual risks arising from the decision.

The combination of the probability and the value of a positive outcome is the opportunity associated with that outcome, given the decision. So:

== there is an individual opportunity associated with each positive outcome and its probability given the decision, and

== the total opportunity arising from the decision, provided the outcomes are independent of one another, is the combination of all the individual opportunities arising from the decision.

Sidebar 2: The Value of Outcomes

The distinction between extrinsic and intrinsic value is very old. Things are extrinsically valued when they are valued for the contribution they make to other things that are valued. Things are intrinsically valued when they are valued in themselves. When I refer to the value of outcomes I mean their intrinsic value.

In Figure 1 I've assumed that you can measure the intensity of intrinsic value on a numerical scale, since I've expressed all the individual values of the 'V's as decimals. I've also assumed that intrinsic value intensity can be expressed as either negative or positive depending on whether the value is viewed as good or bad, or whether you have an obligation to bring it about or to prevent its occurrence. Neither of these assumptions is problematic because techniques have been developed in the field of Decision Theory to produce such measurements.

The idea of 'combination' is critical to the way I've formulated total risk and opportunity. But there's no one necessary way to formalize this idea. Different rules of combination apply depending on how probability and value, both negative and positive, are measured. The most common interpretation of 'combination' is based on measurement methods that produce real numbers for both probability and value. In that case, probability and value are multiplied to arrive at risk and opportunity, (as in Figure 1) and individual risks and opportunities are each added to arrive at total risk and total opportunity. But this is not the only way of combining probability and value and defining risk. So keep in mind that the broader and vaguer perspective can lead to alternative models of risk when the need for them arises.

One way to look at this view of 'risk' is that when we make a decision, we have a kind of 'model' in mind that connects the decision to its possible outcomes, and also includes ideas about the probability and value of each of the outcomes. However, most of us, most of the time, just have some interconnected vague ideas and predispositions that are the beginning of such a "mental model". Then, when we have a tough decision to make, some of us may think much more carefully about the decision and its implications, and even take the trouble to write down our ideas about the outcomes we expect, the probability of each, and the values we attach to them. If the decision being made is important enough, we can even work with others to develop a formal decision model to make our ideas about risks and opportunities completely explicit. But this kind of formality is only worth it if the decision is about matters of great value. Also, please notice that these decision models we formulate are just our ideas about the real world effects of our decisions. They may or may not reflect how the world really works.

So, what happens if our decision model is wrong? What happens to the risks and opportunities we've envisioned?

This brings us to the second category of risk. This is the risk we take of making an error, when we accept that a particular decision model is true when, in fact, it is false.

- The risk of an error in selecting a decision model, and therefore a decision, is "the probability" that you've made the error, combined with the negative consequences of it.

But these consequences may be specified only partly, or not at all, in the model you've selected, because the real world may correspond to an alternative model that you may or may not have formulated, giving rise to its probabilities and its negative consequences. In other words, its risks, rather than those envisioned in your chosen model.

Put a bit differently, any decision has some consequences that its decision maker anticipates. Some of these are positive, and some are negative. Either way, the risk being taken by the decision maker is that whether or not they occur, other, unanticipated, negative consequences, envisioned in the true decision model, will occur, either instead of, or in addition to, the consequences the decision maker anticipates.

So when we commit to a decision and an associated model of opportunities and risks, we take the risk that our decision will not "buy" us the expectations given in our model; but, instead will "buy" us another set of outcomes specified in one of the decision models we rejected. Or perhaps that might have been specified in a possible model we haven't even thought of.

And the risks existing in the world and specified in that alternative model may be either much greater or much less than the risks specified in the model we chose to accept. We have no way of knowing whether this even may possibly be true, unless the alternative model happens to be one we have thought of, and mistakenly rejected. What we do know is that the greater the likelihood that we have erred in our decision model, the greater the likelihood that we will be ignorant about the risks we are taking when we accept it, in preference to its alternatives.

In this formulation, it's very important to keep in mind that the idea of the risk of error is tied to the idea of choice among alternative decision models and decisions, and also the idea that the risk of error involves the negative consequences that you don't think will occur or that you haven't thought of, and not the negative consequences that you think you know will occur as a result of your decision.

Here are a few more thoughts on the risk of error in one's decision model. First, if an error in decision making is selecting and implementing a decision that doesn't "produce" the consequences you expect, but instead produces other consequences, then, all errors are not bad. Sometimes a decision produces positive consequences you haven't anticipated, without producing either the negative consequences you expected, or other unanticipated negative

consequences. So, sometimes an error-filled decision is an opportunity, rather than a risk.

Second, 'probability' is an ambiguous term, isn't it? Earlier, I used it to refer to the propensity of a decision to produce a particular outcome given the decision's context. But, when I refer to the "probability" of a decision model, I have in mind, not the probability of formulating the decision model, or of the decision model's consequences occurring, but, instead, something like: its likelihood of being true in light of the results of attempts to test, criticize, and evaluate it.

This idea of 'probability' is common in ordinary language where we say that one model is 'probably' true while another is 'probably' false because the first has survived the results of our tests, while the second has not. However, this idea of 'probability', it turns out, is not a 'probability' in a sense of that term that is logically consistent with the axiomatic system of probability, sometimes called the probability calculus. So, it is not a 'probability' in the same sense as the 'relative frequency', or 'propensity', or 'set-theoretical', or 'logical' ideas of 'probability' are, since all of these provide consistent interpretations of the calculus. For this reason, I won't call the 'probability' of truth of a decision model with respect to the results of tests, criticisms, and evaluations, a 'probability', but instead will refer to it as 'the truthlikeness' of a decision model and will view it as a performance index for comparing decision models.

Third, considering that risk is probability multiplied by the severity of a negative outcome, which 'probabilities' and which negative outcomes should we take account of in estimating the risk of error? Here are some rules for computing the estimated risk of error assuming that all decision models significant for taking action are evaluated. The rules also specify the idea of the risk of error in one's decision models more clearly.

1. Don't use the negative outcomes and probabilities envisioned in the decision model you accept to assess the risk of error. These are your best guesses related to the first category of risk, the real risk arising out of your decision. But they have nothing to do with the risk of error you're taking by accepting your decision model. Because error in your model means that the negative consequences in it won't occur unless they're also specified in a true alternative model

2. Do use the negative outcomes in alternative models

3. Use the probabilities (propensities) in alternative models along with the negative outcomes, to compute the total risk (of the first type) specified in each alternative model

4. Compute the truthlikeness of each alternative model. Your results are the 'probabilities' (in one of the ordinary language senses) of the various models. The higher the truthlikeness of an alternative, the higher the risk in rejecting it, other things being equal

5. Weight the total real risk specified in each alternative by the truthlikeness of that model. The result is the adjusted total risk of error in each model

6. Assuming that one of the alternative decision models is true, compute the average adjusted risk over all other alternative models

What if none of the decision models is true, including the one we accept as the best one? How can we estimate the risk of error then? Also, how can we go through the kind of risk of error estimate outlined above in practical situations, since there's generally no time for this kind of formal consideration when making most decisions.

Figure 1.2: Average Adjusted Total Risk of Error

The only way to be absolutely right about the risk of error is to know which alternative to the selected model is true and to compute the total risk in that model. But we can't know that, so our best estimate has to be an average of the weighted risks in all the important alternative models. This estimate will be good, however, only if the set of alternative decision models being considered are 'representative' of the range of relevant models. When I discuss how to reduce risks later on I'll explain that representative 'fair comparison' sets are necessary for "killing your worst ideas."

Is it generally practical to estimate the risk of error as I've just indicated in business situations? Actually, no, and I haven't outlined my formal approach to estimating the risk of error in a decision model, because I don't expect you to use it as a matter of general practice.

What I'm after, instead, is your understanding of what it takes to provide comprehensive risk estimation, because once you understand that, you'll (a) know what can be done in those rare situations where the time, importance, resources, and potential negative consequences associated with a decision, merit very careful and comprehensive estimation of the risk of error in accepting a particular model, and (b) be in a much better position to understand what I will say later on about how to reduce risk in our decisions.

Lastly, I didn't provide the formula for truthlikeness because it would have distracted from the main ideas I want to deal with here, immerse us in technical details, and stand in the way of an understanding of my overall perspective.

The Connection to Knowledge Processing and Knowledge Management

I'd like to summarize where we are now, before moving on to how we may reduce risk by killing our worst ideas. I've talked about two categories of risk. The first category is risk that results from our decisions and their propensities to produce negative outcomes. While talking about this kind of risk, I also mentioned, in passing, the opportunities that result from the propensities of our decisions to have positive consequences.

The second category of risks arises from our acceptance of a decision model. It is the risk of error in accepting a decision model. In explaining this sort of risk, I developed the idea that the risk of error in accepting a decision model is the average of the weighted (by 'truthlikeness') total risk of each alternative model in a fair comparison set.

The first category of decision risk is real. It is the real risk we take, whenever we make a decision, of realizing negative consequences. But the second category of risk, while no less real, is our risk of realizing a particular negative consequence, the consequence of being in error, of being wrong, about the risks we are taking. And if this risk is realized and we are wrong, then we may be exposed to risks that we have never thought of, to risks that are greater than any we have imagined.

And while it may be true that an error in the decision model we have accepted also can bring us opportunities we have never thought of, business management is about planning for success, realizing envisioned opportunities, and avoiding risks that are so great that they threaten the very existence of our businesses. It is not about just being lucky, though we all know that "luck" plays its part too.

The problem with the risk of error is that it exposes us to unknown risks, to risks that we do not prepare for, and perhaps to risks we cannot survive. But even more fundamentally, I think that risk reduction must always be about reducing the risk of error in our decision models, because only if we do that can we be confident that decisions that we believe are only moderately risky, will not, in fact, bring us the uncomfortable surprise of unforeseen risks and eventually costs.

We cannot reduce the first category of risk directly, and on a sustained basis, without first reducing the risk of error in our decision models, since they determine what our decisions will be and how risk-filled they are. And like it or not, we cannot reduce the risk of error in our decision models without improving the quality of the decision models we accept, and indeed the general quality of the decision models that we consider, compare, and select among. Reducing risk, in other words, is about making better knowledge, since our decision models are as much knowledge, as any other mental or explicit models we may formulate.

I'm not saying that if we increase the quality of our general knowledge we will reduce risk. I am saying, instead, that if we increase the quality of the specific decision models we use in key domains where there is a lot at stake, we will reduce our risks and improve our outcomes. Now, let's move on to how we can make higher quality decision knowledge in pursuing our efforts to reduce risk.

In the area of making new knowledge, my conceptual framework (originating with Karl Popper, and developed using terms like knowledge claim formulation and knowledge claim evaluation in other publications) says that there are three essential steps: (1) seeking, recognizing, and formulating problems with current decision models, (2) creating new alternative solutions (decision models), and (3) killing your worst ideas.

Improving the quality of decision models and reducing the risk of error means improving performance in each of these areas and in each important domain of risky decision making in your organization. Ultimately, whether improvement happens is dependent on your knowledge workers and on how well they solve problems, but Knowledge Management can enhance each of these three steps by affecting the ecology of knowledge making providing the context for each of these steps. I don't have the space here to lay out even in general terms how KM initiatives might do this for steps (1) and (2). But the rest of this paper will provide a general idea of how it may be done for "killing your worst ideas."

Killing Your Worst Ideas

There's a difference between routine and creative learning. During routine learning from experience we operate according to previously internalized rules that we apply to specific situations. For example, we have general previous mental knowledge about how to make a right turn to steer our cars around corners. Exactly how we steer is affected by the specific characteristics of the road, the corner, how our car may be functioning, and other conditions and circumstances that we learn about through perceiving and thinking about the situation. This routine learning involves both steering trials and, when we receive feedback from the situation about the results of our attempt to steer around the corner, error corrections.

When we turn corners, one corner of our mind, however quickly, plans our turn, we then act to steer around the corner, and we then both monitor and evaluate the results of our steering, all very quickly. If our turn, as seen through our monitoring and evaluating efforts, is not yielding the desired result, if it appears too wide, or too narrow, or too slow, or too fast, we very quickly adjust to our routine learning of these facts, by immediately changing our intent and our steering a bit, and then again monitor and evaluate the result.

The important thing about this kind of routine learning cycle is that it is a trial and error cycle. If we've made an error in our original decision, we can't correct it without first receiving feedback about its consequences. Even in the case of steering a car around a corner, we must experience the effects of our initial decisions about steering before we can do error corrections, assuming they are needed. So, in routine learning, we need to experience the effect of our decisions and the ideas reflected in them, before we can do anything about the errors, if any, in them.

Metaphorically speaking, but sometimes not so metaphorically, in pure routine experiential learning, we cannot "kill our worst ideas before they kill us", but must experience the consequences of our worst ideas, both "trial and error", before we can benefit from our experience. Learning from experience is, too often, not learning the easy way, but learning the hard way; from the sanctions we receive from the real world. It is a process that sometimes forces us to perish along with our bad ideas, rather than having them "die in our stead."

But, don't we all learn better by experiencing the results of our actions and by learning from our mistakes? I think we, along with other animals, do learn best from experience and also from our mistakes, so I'm certainly not saying that we ought not to learn from experience. But evolution has also equipped us with a creative learning process in which criticism, testing, and evaluation can provide a less risky form of pre-instrumental experience in helping us to decide, frequently, but not always with success, which of our decision models are wrong and which may be true. Rather than being a trial and error learning process, creative learning, when done right, is a trial and error-elimination learning process.

Creative learning is dependent on human language. It co-evolved with language and our ability to use it both to create alternative solutions (in this case decision models) and to select among them by criticizing them and evaluating them in other ways. It is our ability, using language, to externalize and objectify our solutions, and expose them to criticism without actually performing them, that makes creative learning the trial and error-elimination-based learning process that it is.

Ultimately, the results of creative learning must also be exposed to experience in practical action, and then, if our creative learning produced the wrong answer, we can still learn from our mistakes, if we have to. But if we're fortunate, creative learning will produce decision models that both survive our tests and criticisms, and also work, and then we will have avoided the greater risk inherent in using models that haven't been evaluated during creative learning.

Now let's go more deeply into the step within creative learning of killing your worst ideas. I've indicated previously, though not discussed in detail here, that new knowledge making or creative learning begins with problems, and then proceeds deliberately with attempts to acquire or create solutions to them. The alternative decision models produced during creative learning are not immediately used in practical action and exposed to experience. Instead, they're evaluated in the step I've called killing your worst ideas.

While the step of creating alternative models is positive in character, killing your worst ideas is 'negative' in its orientation. Its purpose is to test alternative decision models as severely as possible and to try our best to refute each of them. Our objective in doing this is to ensure that the alternative that survives comparison is the best decision model we've been able to create, and the one it makes sense to rely on when we have to decide.

So, how do you go about killing your worst ideas? You do it through fair critical comparison, an idea I briefly mentioned earlier. The basic idea is one of openness in the sense of transparency and inclusiveness, so that the process of fair critical comparison is distributed among all who are interested in participating in solving a particular problem.

Fair critical comparison is a distinctly neo-Darwinian process that focuses on comparing and selecting among competing ideas, and self-organization of distributed critical activities, around the problems being addressed by creative learning. It is the way you should apply criticism and testing in both killing your worst ideas and severely testing the strength of your best ones.

- The first step in fair critical comparison is to satisfy conditions that are necessary for fairness.
- The second step is to compare the members of the comparison set on some set of criteria applied equally to each member.

I'll tell you about the necessary conditions first. There are four of them.

The first is that competing ideas be developed with equal specificity before you compare them. I call this equivalent specification. You can see the importance of this. Laying one alternative out in detail and another in a vague manner could bias an evaluation either for or against one of the alternatives. To ensure fairness in the evaluation you need to be equally specific in comparing all of them.

Second, when preparing competing ideas for comparison and specifying them in the process, more concrete versions of the various solutions must not change the fundamental ideas expressed in their original, less specific versions. This is the norm of continuity in fair critical comparison. Without it, the comparison doesn't really compare what it sets out to compare.

Third, often, when you set out to compare different ideas or solutions you may find that they are expressed in terminology that's so different that the solutions defy easy fair critical comparison. In this situation, you need to create a broader framework of concepts that allows you to fairly compare alternatives. This is the norm of ensuring commensurability in fair critical comparison.

Fourth, there's nothing so unfair to an alternative than to exclude it from a comparison. A solution can't become the solution, if it doesn't even get compared to others. You've got to do the best you can to ensure that comparison sets of solutions are representative of the range of available solutions, and this is one reason why I strongly recommend distributed participation in developing the alternatives. In killing your worst ideas you need to make sure that the fair critical comparison set you'll be using is as complete and representative as possible. This is the norm of completeness of the comparison set.

Viewing these four requirements for fair comparison, two questions immediately arise. One, how do we know these necessary conditions are all that we need for "fairness?" And two, why does the comparison among alternative solutions have to be "critical"?

To answer the first question, I don't know whether my four necessary conditions are all you need for fair comparison. They're the ones I've been able to come up with in some years of thinking about fair comparison, but I certainly invite you to add to these if you can think of others that are independent of the four.

Getting to the second question, there are a couple of reasons why fair comparison needs to be critical. (1) Any comparison ordering solutions according to some normative criterion is implicitly critical. If you're comparing two solutions in terms of their simplicity, both will fall short of perfect simplicity. The one that is farthest from the ideal of simplicity has a lower rating or ranking than the second solution, so the result of the comparison is implicitly critical of the second solution.

(2) If the decision models you're comparing are based on generalizations, and you're trying to choose between two of these models by looking at data, you need to find data that refutes one of them. Data that supports both won't help you choose. So, when you test your solutions, and they are based on generalizations, those tests have to be critical if you want to select a solution.

The bottom line here is that in comparing solutions, you're looking for failure, not success. You're trying to get your alternatives to fail during your evaluation within the creative learning process, so you'll know which solution, or possibly solutions, would fail you if you applied them in practical action. Again, you're operating critically, and yes, negatively, because you're trying to kill your worst ideas before they can hurt you.

An important objection to this idea is embodied in the following line of questioning. If killing our worst ideas involves a critical approach we can't avoid, won't it be difficult for us to do it well on a continuing basis because of the natural aversion of people to criticism? What sort of critical style can we use to be successful at fair critical comparison? Won't this approach result in uncontrollable conflict?

It's not easy to apply the approach of fair critical comparison on a sustainable basis, without considering issues such as how we overcome aversion to criticism, how we moderate conflict in the face of criticism, and the sort of critical style we ought to use in performing comparisons. We'll cover the first two issues as we move along in our discussion below. As for the kind of critical style that should be used, it must be very disciplined and civil.

What I mean by this is that rhetorical devices in argument including • ad hominem attacks on people who favor one or another competing solution, • labeling of solutions you don't like using terminology with negative connotations, • personal attacks on those who favor solutions to problems you oppose, and • appeals to authority,

must be off-limits in performing fair critical comparison because they create anger and stress without contributing to fair critical comparison in any constructive way. The style of criticism used must be one of cooperative exchange in a common effort to arrive at the best solution.

The thing to keep in mind here is that fair critical comparison is about providing the toughest tests you can for all alternatives so you can expose the ones that won't work, and, hopefully leave some survivors standing. It is not about defending any particular alternative, even if you originated it. You won't get credit for a solution that survives criticism because you unfairly, and uncivilly, defended it, but will only fail when it is tried in practice. So what's the point of defending a point of view by using ad hominems, labeling, or the other devices I just mentioned for winning arguments? If you win the argument you'll only lose anyway when your bad solution doesn't work.

I haven't provided specifics about standards in fair critical comparison, so let's go there now. The first thing to remember is that there are no hard and fast critical 'standards.' What exists are frameworks comprised of loose critical perspectives that people use in flexible ways to evaluate competing ideas. The frameworks are subject to change and modification since anyone can come up with novel perspectives for criticizing competing ideas at any time. Here is the framework for critical evaluation I recommend; but remember, it's just a theory.

- Logical consistency (inconsistent decision models are invalid and must be reformulated)
- Empirical fit (competing models fit current and past data to varying degrees)
- Projectibility (models vary in their plausibility, models

vary in their after the fact success in prediction) • Systematic fruitfulness (extent to which a decision model facilitates novel deductions) • Heuristic quality (extent to which a decision model facilitates new conjectures) • Systematic coherence (coherence of statements relating abstractions in decision models, coherence of statements relating abstractions and concrete terms in decision models) • Simplicity (economy in number of variables in a decision model, simplicity of mathematical form) • Estimated risk of error in accepting a model.

As we saw earlier, the idea of the risk of error requires that we take value assessments and probabilities into account in estimating risk as well as "truthlikeness." But since when do value assessments enter the process of deciding which of a number of competing models is closer to being true? Isn't this kind of selection process supposed to be objective? How can we both incorporate value assessments and still be objective?

How can we not and still be objective? Let's imagine three solutions to a problem and let's say we compare them on all other criteria except the risk of error. Let's also imagine that solution #1 looks like it's the best (most "truthlike") of the three on all the criteria we've used for the comparison. Solution #2 hasn't performed as well as solution #1, but it hasn't been falsified by data, or at least tests that appear to falsify it have been met with serious criticisms about their validity. Finally solution #3 has been clearly falsified by tests we've run.

So far it looks like objectivity demands that we rely on solution #1 in making a decision. However, the decision model suggesting solution #1 predicts that if we take action 'a' we'll get a better result than if we take action 'b', but that action 'b' produces an acceptable result.

On the other hand, solution #2 predicts the same result as solution #1 for action 'b', but predicts that the result of action 'a' will be disastrous in producing a very negative outcome. Now, is it better to rely on solution #1 in such a situation rather than solution #2? And whether or not it is, are we "less objective" if we allow the prediction of a highly negative result by the decision model underlying solution #2 to affect our decision about whether we should accept solution #1 as the solution to problem and the basis for our decision?

This example of the three solutions illustrates how the risk of error in accepting a solution, such as solution #1, should affect our decisions to accept a solution that initially seems most "truthlike" compared to competitors. That is, the risk of error should not wholly determine our assessments of truthlikeness, but it should affect our decisions increasingly, as the seriousness or risk of error increases.

But why this should be so? What does the risk of error in a decision model, as represented in alternative models, have to do with whether a model is true or false?

The answer is that we have to distinguish the question of whether a model is true or false, which, in the end depends only on its correspondence with reality, from the question of how we will measure truthlikeness for purposes of deciding what we will provisionally take to be true when we have to make a decision. We have no way of directly comparing our models with reality. All we can do is to compare them with one another in the light of criteria like the ones I just listed. The results of such comparisons may always be wrong because we and our measurements are fallible. Our method of fair critical comparison can kill our worst ideas over time and reduce the risks associated with our errors, but, and this is a big 'but', we need time for fair critical comparison to work, and therefore we need to have a fair critical comparison process that's sustainable enough that it has time to work to eliminate errors.

If our method of measurement of "truthlikeness" lets us accept solutions and decision models that don't take the risk of error into account, the consequence may be that our businesses, and the processes of fair critical comparison within them that we rely on to more closely approach the truth, may be precipitously undermined. In other words, an undue emphasis, in the short run, on epistemic evaluation criteria in comparing competing solutions and decision models, and a failure to take account of the risk of error, may lead to the destruction of our efforts to create more "truthlike" new solutions (knowledge) in the long-run. So, the role of the risk of error in evaluating "truthlikeness" is to moderate the bias toward unsustainability in fair critical comparison, by introducing a hedge against the fallibility of ourselves and our measurement processes.

Having said all this, I now need to say plainly that just as there are no rules that tell us for certain which criteria for fair comparison to use in killing our worst ideas, there are also no such rules for telling us how heavily we should weight the risk of error in deciding whether or not to accept an epistemically superior decision model in preference to one that has not performed as well.

Remember when we talked about the impracticality of estimating the risk of error in decision models using a formal methodology? I pointed out that for most decisions the effort wasn't worth the cost, and I also pointed out that for cases where it was worth the cost I could specify such a methodology. The same applies here to "killing your worst ideas." As a matter of general practice you won't want to use formal methods with clear rules, but will want to rely on informal exchanges among people who want to participate in solving a problem, and who will create their own criteria on the fly, as it were, for comparing decision models.

When the stakes are high, however, then you'll want to develop formal decision models for doing fair critical comparison. Methodology, methods, techniques, and tools you can use to develop your own formal models for fair critical comparison reflecting the priorities of your own organization, are largely beyond the scope of this brief book. However, Sidebar 3 offers some references and guidelines for implementing a formal approach to fair critical

comparison.

Sidebar 3: Formal Approaches to Fair Critical Comparison: References and Guidelines

You'll find two formal approaches to measuring truthlikeness in the Appendix to Chapter 5 of my book with Mark McElroy, *Key Issues in the New Knowledge Management*, Burlington, MA: KMC Press/Butterworth-Heinemann, 2003. There I present an approach based on the Analytic Hierarchy Process modeling technique and a second based on fuzzy measurement modeling. Both approaches should be modified in line with the account of risk of error I've given earlier.

You can accomplish this, by first computing estimates of 'truthlikeness' for each competing model in a comparison set and then subtracting from each of these estimates the average adjusted risk of error of all models alternative to the one whose 'truthlikeness' is being estimated. The resulting set of adjusted estimates will not be the final measured values of the 'truthlikeness' of the competing models but will be an intermediate set of estimates.

Using the new adjusted estimates, the procedure should then be repeated until each of the 'truthlikeness' estimates converges toward a final value. The resulting set of converged values will be a measure of the truthlikeness of each alternative model incorporating the estimated risk of error criterion mentioned earlier.

"Killing your worst ideas" is part of a broader process of open, distributed problem solving characterized by transparency, inclusiveness, and trust. So, there is no time when criticism should end. Criticism should be a continuing process. Even when we accept a solution as provisionally true, or at least as the most "truthlike" among our alternatives, we should continue, even while acting in accordance with the solution, to continue to surface problems with it, in an effort to ensure that it is a good solution that can continue to survive all that both we and reality can throw at it.

The key to distinguishing cultural knowledge from information, and to the growth of knowledge through "killing your worst ideas" is your record of how competing decision models have performed in relation to each other, and your criteria for comparison. So tracking results is of fundamental importance to making new knowledge and reducing the risk of error in your decision models. Information technology applications can help with tracking and creating the performance record you're seeking. But for reasons of space I can't offer a discussion here.

What group environments are effective in "killing our worst ideas?" Pretty much the same kinds of group environments as are important in generating new models. Whatever group environments are used must be open to fair critical comparison and killing our worst ideas through criticism. That means transparency, inclusiveness, and trust must be characteristic of these environments. In particular, collaborative teams, facilitative groups, and communities of inquiry can all be very important in fair critical comparison.

Criticism, like other parts of the problem solving process, should be distributed and open. But, won't distributed criticism undermine management authority? In the beginning of "openness," fair critical comparison might seem disrespectful of authority. But, imagine an enterprise in which "openness" is the order of the day. Everyone is subject to fair critical examination of their decision models. Criticism in such an environment, if it is not personal, only serves as a social support to ensure that surviving models are the strongest we have. The result is that everyone's decisions are better and authority is strengthened, not undermined.

During the transition to "openness", everyone needs to be aware that competing decision models will be questioned and that most will be found wanting. That sort of result is the purpose of instituting a policy of openness toward fair critical comparison. It is the result management wants, so there's no reason it should be taken as undermining its authority.

If Management's favored solutions don't do so well in fair critical comparison, then it has the choice of acting according to its favored solution, or other solutions that have performed better during the knowledge making process. Management's authority and responsibility for making decisions is entirely intact. Higher quality knowledge making just enhances Management's options and potentially its performance.

Communities of inquiry can be very important in killing our worst ideas, provided they are true communities of inquiry, are open in their fair critical comparison processes, and are characterized by civility rules such as those listed earlier. If collaborative teams have open critical processes they'll also be very useful in killing our worst ideas. Otherwise, they won't work well for fair critical comparison.

Historically, Group Decision Making techniques such as Delphi, Nominal Group Technique, Joint Application Design sessions, Team Analytical Hierarchy Process, and facilitation techniques in general have been directed toward finding the consensus of a group. But the key to making Group Decision Making techniques work for fair critical comparison is to orient them toward criticism and managed conflict, rather than toward achieving consensus.

Group Value Measurement Technique (GVMT), one I've had considerable experience with, was often used to develop alternative models and to encourage critical evaluation of model alternatives. It emphasized statistical analysis of judgments and ratings in a manner designed to identify variety rather than consensus. Knowledge Cafés also work in such a way that alternative views emerge from the process.

Analytical modeling can be very important when it's used to measure empirical fit, forecast success, and to set up tests by generating the logical consequences of alternative models in simulations. Analytical modeling techniques can also be very useful when we're using formal models for performing fair critical comparison.

There are a wide variety of other contexts for fair critical comparison. They include personal networking, project meetings, presentations of various kinds, story-telling, and any other context where people meet and exchange views. In each of these contexts the issue for fair critical comparison is the same. Are the contexts open? Are they transparent, inclusive, and characterized by at least moderate trust levels? If so, they can be important in contributing to fair critical comparison.

Project meetings are one of the most important contexts for problem solving in business organizations. But very little attention is given to how they ought to work as environments for evaluating alternative solutions. Introducing openness and fair critical comparison into project meetings may be difficult. But getting it introduced throughout your organization is essential, if you really want to do a good job of killing your worst ideas.

Recently, story-telling has become quite the rage in corporate circles and there is no doubt that stories are a very effective method of presenting ideas and helping people to feel empathy with the case being presented in a story. Stories are particularly effective when you want to sell your point of view to busy people who want to be given the essence of your ideas in a brief, interesting and attractive form. But a great philosopher once pointed out that criticism and argument were developed by humans to defend us from stories and story-telling, and specifically to expose their less than truthful elements.

From my perspective, creating stories corresponds to generating and presenting new ideas. And the best way to use stories in making new knowledge is to create competing stories, and then use fair critical comparison to select among them. In other words, I think that openness and fair critical comparison are just as important in storytelling as in any other context. However, current works on storytelling don't emphasize openness and fair critical comparison, but instead emphasize how to tell effective stories to present one's point of view. In other words, they stop the process of knowledge making at creating new ideas, but instead of attempting to kill our worst ideas, they attempt to justify new ideas by making their presentation as persuasive as possible.

In the area of outcome metrics, I think the ratio of relevant surviving knowledge claims to the problems they are addressing is very important as is the ratio of effective surviving knowledge claims to problems. In addition, the average cycle time of fair critical comparison is important. I can also specify a host of process metrics to help in causal modeling. But the "bottom line" metrics are the ones I've listed above. In other words, the key metrics are: relevance, impact, and efficiency.

People will participate in the distributed critical process if they think the problem to be solved is somehow their problem, or alternatively if the challenge represented by the problem engages them. In most situations this won't be many people, but it will still be many more than participate in closed evaluation processes, and that will make all the difference for the quality of fair critical comparison, and for your success in killing your worst ideas.

From an organizational point of view, overcoming aversion to criticism can be done by implementing a long-term, multi-faceted program to replace the knowledge that criticism has too many bad effects and should be avoided in favor of "positive thinking," with the knowledge that criticism is necessary if we're going to kill our worst ideas. In a company where top management will support the transition you can:

- make policies of openness in eliminating errors in knowledge claims through fair critical comparison,
- have training workshops for managers and knowledge workers in critical thinking, critical styles, innovation, and fair critical comparison in team and project environments,
- establish communities of inquiry dedicated to problem solving through killing your worst ideas,
- employ Group Decision Processes that emphasize critical approaches,
- hire people with critical skills and favorable attitudes toward critical approaches, and
- modify business processes, so that criticism is introduced into the decision making process in an impersonal way that facilitates creative learning by both the individual and the organization.

The last item on this list is particularly important because it can be used to gradually introduce the critical outlook into organizations, even when top management doesn't want to establish an explicit policy of introducing fair critical comparison during error elimination.

Some of the training I mentioned above can help people to use criticism and be more comfortable with it, as can communities of inquiry. But information systems that can integrate people into the flow of critical ideas in a natural way related to their job roles, can also be very effective. While they do their jobs, people can receive and give criticism in an impersonal way that is easy for them to live with.

Moderating the inevitable conflict resulting from critical exchanges depends very much on training people in criticism and on enforcing civility rules. Managers and knowledge workers involved in critical exchanges must not allow these to get personal or political. I know it's hard to prevent this, but if norms for critical exchange are supported by everyone, things will get personal less often and conflicts can be moderated.

So, in general terms, these are my views on how to reduce the risk of error, how to make new knowledge, how to seek, recognize and formulate problems, how to generate new ideas, and how to kill your worst ideas. The result of killing your worst ideas or, more formally, fair critical comparison, is, even when you're doing it informally, at least a

classification of alternative decision models into three categories: the surviving models (the truthlike models), the models you're undecided about (the less truthlike models), and the models that you think are no good (the false or refuted models). The ones you can most likely rely on are the truthlike models, and the better the job you've done in performing fair critical comparison and killing your worst ideas, the higher the quality of your truthlike models, and the lower your risk of error, will be.