An Action Model for Risk and Uncertainty in Decision Making - How to Avoid the Paralysis of Analysis

Jeffrey A. Robinson, Ph.D.
jeff.robinson@aqionline.com

Abstract

It is fairly well accepted that 'risk' and 'uncertainty' impact decision making activities and make them more complex and difficult. Unfortunately, it is not always clear, from the way we use these terms, how risk and uncertainty influence how we should act when complex decisions are being considered. Part of the problem is that there are several different terms used to refer to these complicating factors. The words 'risk', 'impact' and 'cost' are used almost interchangeably… as are 'probability', 'frequency', 'likelihood', 'certainty', or 'uncertainty'. To a great degree, when any of these factors are present we tend to analyze problems more carefully in order to make the best choices and to avoid potentially undesirable outcomes. There are many different analytical techniques that can used to break down details of a problem, explore the underlying factors to project likely outcomes and prioritize possible choices and courses of action. These include Failure Mode Effects Analysis, regression analysis, correlation analysis, root cause analysis, Ishikawa diagrams, Cause-Effect matrices, and a myriad of quantitative risk analysis techniques. However, this general tendency to analyze things when risk is perceived can be detrimental and can lead to a phenomenon known as the "paralysis of analysis" in which far more time is spent perusing and processing data than is necessary or appropriate. To understand how risk or uncertainty affect the decision making process, we must first carefully define these terms and examine the nature of choice and action taking with a more formal approach. This paper presents two important concepts. The first is a rigorous definition of risk and uncertainty. The second is a taxonomy that clearly suggested different modalities of action for the different combinations of these two factors. Hypothetical case studies are provided to illustrate the rationale behind this action model, which clarifies when quantitative analysis 'is' and 'is not' appropriate. This model can be applied to projects of all types including software development and process improvement projects. If used during project selection and prioritization, it can be effective in avoiding excessive analysis activities and minimize the tendency to overanalyze whenever risk or uncertainty are present in any proportion. Developing a formal methodology for assessing and managing risk is important to optimize utilization of resources and mitigate the effects of potentially undesirable outcomes. Knowing when and how to analyze risk is important, but it is equally important to understand when such analysis is not warranted and would only cause unnecessary cost and delay. Managers and decision makers need to understand that decisive action without excessive risk analysis is still possible when risk and uncertainty occur alone. Only then will they be able to act confidently and appropriately when either or both are present.

Biography

Dr. Jeffrey Robinson is the Vice President, Principal Consultant, and Co-founder of Accelerated Quality Improvement, a leader in process improvement consulting and training. He is the co-author of Six Sigma Software Quality Improvement – Success Stories from the High-Tech Industry, McGraw Hill, 2011. Dr. Robinson is a senior technologist and a Six Sigma consultant with more than 25 years experience in IT and manufacturing with companies such as Motorola, Medtronics, and Rockwell. He has served as Master Black Belt, Division Quality Manager, Manager of CIM (Computer Integrated Manufacturing), and System Manager of large automated factories. He is a frequent author and lecturer and has taught many Six Sigma courses.

1 Introduction

It is fairly well accepted that 'risk' and 'uncertainty' impact decision making activities and make them more complex and difficult. (Riabackel, 2006). Unfortunately, because of the way we use these terms, it is not always clear how we should act when complex decisions are being considered. (Abdellaoui and Hey 2008)

Part of the problem with the way that risk and uncertainty affect decision making is that there are several different terms used to refer to these complicating factors. The words 'risk', 'impact', and 'cost' are used almost interchangeably… as are 'probability', 'frequency', 'likelihood', 'certainty', or 'uncertainty'.

To a great degree, when either of these factors are present in any measure, we tend to analyze problems more carefully to help us make the best choices and to avoid potentially undesirable outcomes. The discussions that follow focus on the choices and actions that people make when there is uncertainty about what to do or how to act. Sometimes, what people decide to do is to conduct further analysis, to help them make a more accurate choice or selection. Sometimes people act directly, choosing a specific alternative that they have identified or a solution or action plan. The dilemma that decision makers face lies in how risk and uncertainty combine to make this selection of appropriate actions more difficult.

There are many different analytical techniques that can used to break down details of a problem, explore the underlying factors to project likely outcomes, and prioritize possible choices and courses of action. These include Failure Mode Effects Analysis, regression analysis, correlation analysis, root cause analysis, Ishikawa diagrams, Cause-Effect matrices, and a myriad of quantitative risk analysis techniques. However, this general tendency to analyze things can be detrimental and can lead to a phenomenon known as the "paralysis of analysis" in which more time is spent studying and processing data than is necessary or appropriate. (Bensoussan, Fleisher, 2008)

To understand how risk or uncertainty affect the decision making process and to determine when extensive analysis is appropriate (and when it is not), we must first carefully define these terms and examine the nature of choice and action taking with a more formal approach. (Harrison 1998 and Sewell 2009)

2 Definitions

First, let's examine the differences between Risk and Uncertainty.

Risk maps most closely to COST, that is the cost or impact of a particular problem or alternative. Risk is the magnitude of loss (in time, money, effort or lives) associated with a specific choice or decision. It is how much might be lost (or gained) based upon different actions or alternatives, which might result from a decision that is made, or not made. (It is important to remember that taking no action and not selecting an alternative is always a valid alternative in any problem or decision-making process and may indeed be one of the most costly alternatives in the long run.)

Sometimes risk can be quantified quite precisely. At other times risk is probabilistic in nature. Risk is a factor that linearly corresponds to cost. Sometimes risk includes considerations of likelihood, or even detectability, but in all cases risk is proportional to cost. That is, if for a given likelihood or detectability, if the cost of a particular failure or a course of action is twice as much as another, then the risk is also doubled. Thus risk relates most closely to the impact or cost of a particular choice, alternative, action, or outcome (see Figure 1). When other factors are not clear, it is easiest to equate risk with potential dollar loss. When probability is considered, risk is the probabilistic cost. That is, how much money might one lose if a particular event occurs or a particular choice is made. Basically, the greater the potential financial impact, the greater the risk.

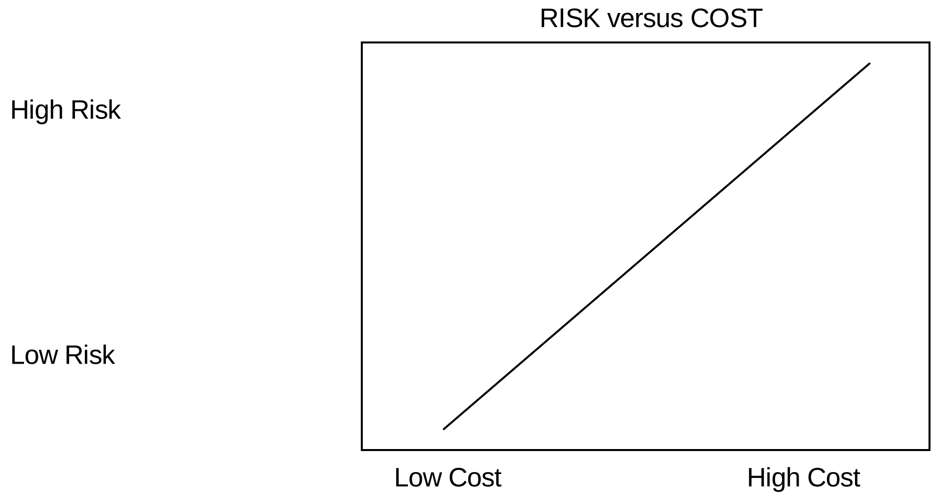

Figure 1 – Risk and Cost

Probability, on the other hand is the likelihood of a given outcome associated with different events, actions or choices. Probability is the degree to which the associated risk (or quantified loss) is likely to occur. It is usually expressed as a number ranging from zero to one. Zero means there is no chance of an event occurring. One is absolute certainty that it will occur. It may also be represented as a percentage from zero to one hundred.

It is important to note, however, that Probability is NOT the same as *Uncertainty*. This is an important distinction, since uncertainty and not risk is a primary drive in selecting appropriate actions and alternatives. Unfortunately probability is too often confused with uncertainty and, while they are related, they are quite distinct from one another. Contrary to what one might initially think, high uncertainty does not occur when probabilities are low.

Indeed, when probabilities are *very high* or *very low* (nearly one or almost zero), there is almost absolute certainty. That is, when the probability of an event is very high, then that thing is almost certain to occur. Conversely when the probability of an event is very low, the event is almost certain to not occur. The greatest uncertainty lies in the range between these two extremes. *Uncertainty* is at its maximum in the middle of this region, when the probability of an event is 50%. (Figure 2)

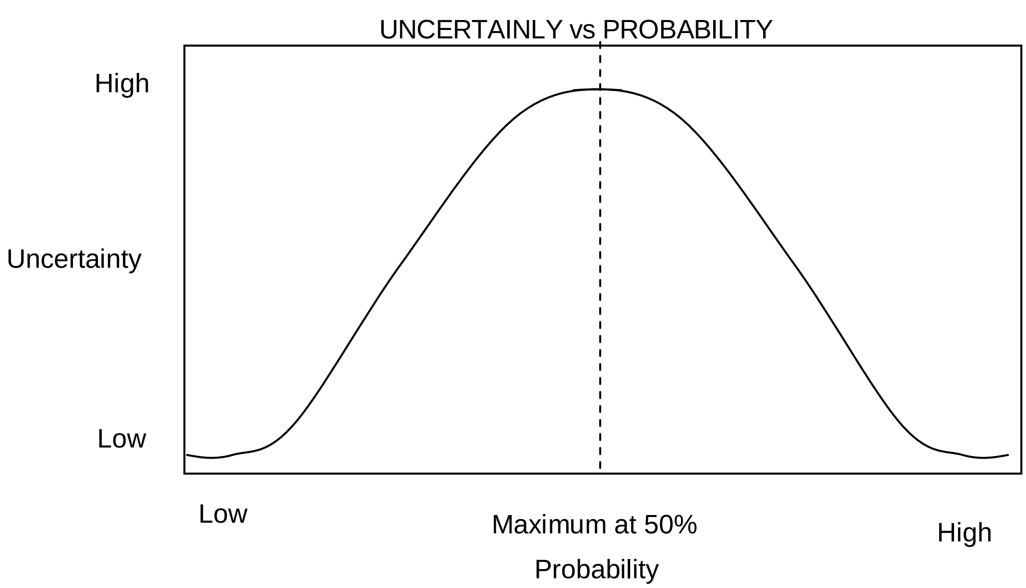

Figure 2 – Uncertainty and Probability

If we consider the situation where probability is very high or low, uncertainty is minimized. Events associated with these extremes are either very certain to occur or to not occur. Uncertainty in these regions is thus very low. It is in the middle that uncertainty is the greatest; when the outcome is as likely to occur as not. (Remember then that uncertainty is not the same as probability.)

3 A Taxonomy of Uncertainty and Risk

Now, let's examine different situations when risk occurs in conjunction with uncertainty. Below is a graphical taxonomy of uncertainty and risk (Figure 3). In the table below, four different combinations of high and low risk and uncertainty are identified.

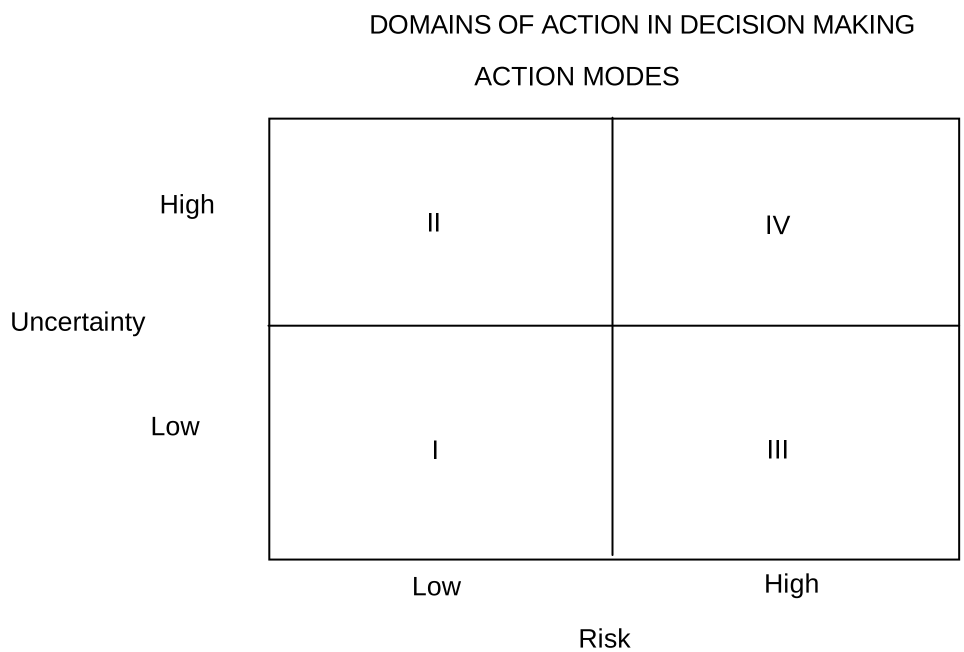

Figure3 – A Taxonomy of Risk and Uncertainty

As we examine these situations we will find that each of these domains (or regions) has a different type of optimum approach for decision-making.

For the sake of discussion each of these regions will be numbered.

Region I - Low Uncertainty/Low risk
Region II - High Uncertainty/Low risk
Region III - Low Uncertainty/High risk
Region IV - High Uncertainty/High risk

Let's examine these different domains individually.

3.1 An Analogy

3.1.1 Region I - Low Uncertainty/Low Risk

For the sake of these discussions, we will create a hypothetical problem that has low uncertainty and low risk.

Let’s say that we have a problem with a large, costly software system. An examination of the system reveals a bug in the code. Moreover, the examination of the code reveals that the root cause is associated with a single subroutine and that the fix is a relatively easy one. In this case, uncertainty is low, because we know exactly what solution to implement and the effort to do so (i.e. the cost and therefore the risk) is low

Assuming that there are no additional difficulties with scheduling of performing the correction, the solution to this problem is simple.

Action: Just do it.

Our course of action is clear. Simply fix the erroneous code, test it and correct the problem. Resolving this problem does not require any special root cause analysis, quantitative risk assessment methodologies, or statistical techniques. You just implement the solution and do it. (Don’t we wish all problems were this simple?)

3.1.2 Region II - High Uncertainty/Low risk

Next, let’s consider a different situation. In this case we will modify our original scenario and increase the uncertainty, but not the risk.

Once again, we have a problem with a large complex software system. However, this time we find that a subroutine call to a sorting utility is malfunctioning and the sorting algorithm needs to be rewritten (Perhaps the current method only works on small files and the application will be sorting very large files instead.) In any case the current method is unacceptable and a different one will have to be applied, and there are a half dozen possible methods that can be used. Unfortunately, in this case, it is not certain which of the alternative methods will work the best. While solutions are known, it is not clear which solution will be adequate and there is no clear consensus by the programming team on which method to use. Some prefer one technique; others argue for alternatives. The problem now is that we don’t know which subroutine to write to replace the inadequate one.

The good news is that method is relatively easy to write and the effort would only take a few hours to code and implement. Thus, uncertainty is high, but risk (or cost) is low.

Even if we narrow down the selection of one or two methods, we still face a problematic decision point. We need to determine which algorithm to use and we need to make sure it works adequately. So how do we decide which solution to implement? How do we pick the correct algorithm?

So what course of action do we take?

Action: Select one of the alternatives and test empirically.

That’s right! When the risk of failure (i.e. the cost of an action or choice) is very low, then one need not worry overmuch about making bad choices. Indeed, making fast choices, even though they are the wrong ones, may be far better than trying to only make the best choice. This is called making a satisfying decision rather than a maximizing one.

In other words, the best course of action in this case is to simply write up code corresponding to the most favored options and alternatives and test their ability to sort large files. If the first subroutine doesn't work, try the next one instead. If that one doesn’t work, try the third. If there is no risk to writing these subroutines and trying them out, then the fastest way to find a solution that works may be to systematically test different alternatives and until one is found that works. Moreover, you do not have to test all possible sort methods, because any one that is sufficient would be adequate. This leads to the rapid identification and selection of an appropriate replacement algorithm with minimal delay and (because of low risk or impact) at minimal cost, as well.

Once again detailed analysis of the different methods is not necessary. Experimentation is faster, simpler and cheaper.

3.1.3 Region III - Low Uncertainty/High Risk

Moving on to yet another situation, we will consider the scenario where we have a new high risk but low uncertainty.

In this third case, our system has again broken and needs to be fixed. This time, however, uncertainty is low. We have found that the current database used by the system no longer works. The freeware DBMS that was originally used to process SQL queries has reached its performance limit and the current number of and size of tables exceeds its capabilities. It has been determined that there are a couple other databases that could be used to replace the original DB. Specs indicate that either of the options could simply be plugged in and used with minimal impact to the system. However, this time the impact or risk is very high because the replacement DB for the system is not free and is, in fact, quite expensive. Our problem now is to decide, should we buy either of the new databases? In this particular scenario, Uncertainty is low (we know that either of the chosen solutions will work and that either will exceed current system requirements), but risk is high (since the solution is very expensive).

So what is our best course of action?

Action: Make a judgment.

When this type of problem occurs, one must determine whether or not the purchase of the solution is worth the cost. To answer that question, one typically needs to perform a simple cost/benefit calculation.

For instance, the decision maker might ask several questions, such as:

- How long would the system have to run before it pays for the cost of the new DB?
- Is there sufficient revenue or service provided to justify the upgrade?
- Would it be cheaper to buy a new system? Or can the upgrade be justified?
- Or should the current system be trimmed down to run in a degraded mode?

This is the scenario when quantitative methods can best be used to make the decision. The numbers will determine whether or not the decision to upgrade or not is the right one. There is, however, little analysis to be performed. The action of choice and selection primary consists of making a quantitative judgment. A cost benefit analysis, or the use of weighted attribute analysis could be used to show which alternative is best.

For instance, if the system is a legacy system that is rarely used, there may simply not be enough potential revenue to justify upgrading the system. On the other hand, if the system is a critical component of a core business process that facilitates financial transactions of more than a half a million dollars every day, then you want that system upgraded and fixed as quickly as possible. The choice is based upon a quantitative judgment focusing on the cost of the different options, but there is minimal analysis involved.

3.1.4 Region IV - High Uncertainty/High Risk

In our final situation, we address our worst case scenario, then risk and uncertainty are both very high.

Once again, our system is not working properly. Performance has degraded. information is being lost and the system keeps crashing. This time it is not clear which part of the system is broken or what needs to be done to fix the system. The programmers who write the system a few years ago left the company recently. No one knows where the manuals are and the system is no longer under vendor maintenance.

The cost of the fixing the system could range from fixing a few critical bugs to overhauling or even replacing the system. Thus both uncertainty and risk are high. The possible options are varied and the best solution, or even an adequate one, is not yet obvious to anyone.

So what do we do?

Action: Analyze; conduct additional research.

In this case, one cannot simply go ahead and start writing code or selecting replacements for specific subsystems. It is also not acceptable to use a trial and error method of employing patches across the system, because the cost of such experimentation could be exorbitantly high and costs of delays could exceed the cost of replacing the entire system. Similarly, we cannot conduct a cost/benefit analysis, because we don't know what part of the system needs to be fixed and therefore we cannot assess the value of a solution.

When this type of problem is faced, either uncertainty or risk must be reduced before a decision can be made. Additional study needs to be done to identify which part has broken or how much it will cost to find out and fix it.

Possible methods of analysis might include: examination of symptoms and failure modes, root cause analysis, fault signature analysis, examination of failure metrics, error reports, and implementation of special diagnostics. Several quantitative methods were noted earlier. In any case, multiple alternatives need to be identified and evaluated.

3.2 The Purpose of Analysis

It is important to note the fundamental purpose and objective of analysis. The reason we need to analyze and study the situation more is either to:

1. Reduce uncertainty (and revert to a Region III action mode) or to
2. Reduce risk (and revert to a Region II action mode).

It is not clear which of these two outcomes may result from analysis that is performed. Eventually you will understand with greater certainty the factors associated with the problem or you may discover cheaper and safer alternatives that will reduce the cost or risk of selecting a course of action. In either case, when either of these are achieved, the decision becomes easier as you transition into decision Regions II or III.

At first glance it seems as though the combination of high risk and high uncertainty actually promotes inaction (at least while some subsequent analysis is performed). However, the amount of analysis (and hence the degree of inaction) is itself limited.

No matter how much analysis you perform, you can never completely eliminate either Uncertainty or Risk (though many people try.) You could spend years mulling over different options without taking any action and you would might never completely eliminate uncertainty and find a solution that is guaranteed to succeed. Too much analysis would be devastating. Not only does the analysis cost time and money, every minute the system is down equates to lost productivity or lost revenue. There is, in fact, a maximum amount of analysis that should be performed. Eventually, you could spend more money analyzing the problem that the cost of completely replacing the system..

Indeed, one can quantitatively calculate the amount of analysis that should be performed. This technique is called "The Price of Perfect Information" and equates to the point where the cost of analysis exceeds the cost of committing to (or selecting) alternative courses of action, even if they are wrong (Bodily, 1998; Evans 2007). If you spend more money than this, by researching or analyzing the problem, then you are spending more than the risk of simply guessing or randomly selecting a solution. Unfortunately, some managers, who are risk averse and who cannot adequately differentiate risk and uncertainty, tend to analyze far too often and too long.

In the case of our model above, if the combined cost of delay and the cost of analysis exceeds the total cost of the system then it would be simpler to just buy a new system (COTS) and install it. The very act of not deciding and performing extensive analysis may mandate a specific loss ($) and become a fiscal certainty. (In such a case, the decision reverts to a Region III decision scenario of high cost and low uncertainty.)

3.3 Action Model Summary

Let's review these modalities of suggested action one final time (See figure 4). Note: each domain has a different decision making strategy. remember, the decisions we are making are related to the selection of alternative courses of action or the selection of specific solutions. Risk relates to the cost of making (or not making) a choice. Uncertainty relates to our confidence that the solution will provide the desired outcome or remedy and reflects the degree of understanding about the problem domain.

- In Region I, you simply implement the obvious solution.
- In Region II you can perform empirical tests to find the right solution by trial and error.
- In Region III, you may use quantitative methods to make formal judgments.
- In Region IV, you need to go through a full analytical decision-making process (setting objectives and requirements, identifying and weighing alternatives, selecting alternatives, etc.).

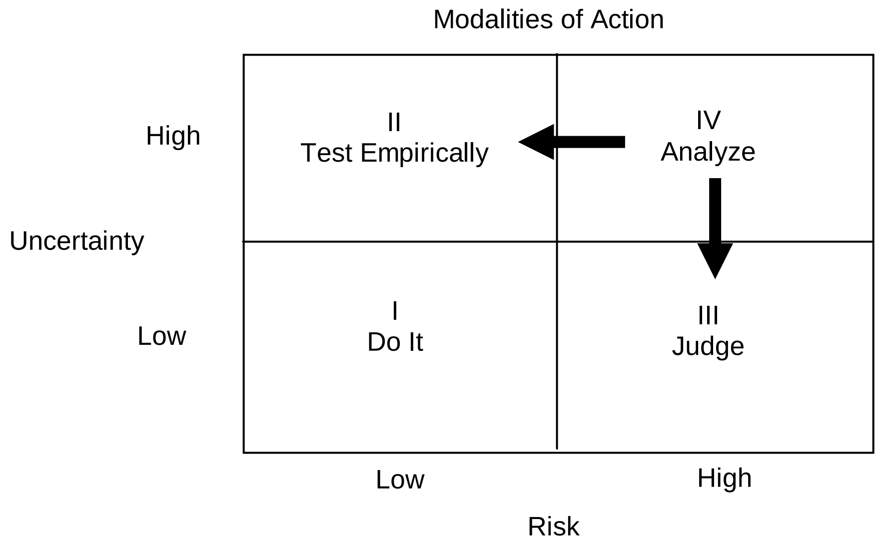

Figure 4 – Objectives of Analysis

In the case of Scenario I (low uncertainty and risk), the decision is an easy one. No extensive analysis is necessary. Similarly, it would be unnecessary to perform a full decision analysis in the case of Scenario II, since it would be faster to test alternatives and try them out than to perform analysis to infer the best choice. Experimenting empirically and testing and failing quickly can lead one to a satisficing decision with far less delay and effort.

Region III also does not warrant full analysis. However, it does justify a cost benefit calculation upon which a decision can be confidently made.

The only domain where extensive analysis is appropriate is Region IV, when both risk and uncertainty are high. Indeed, the purpose of analysis is to reduce uncertainty or risk sufficiently to drive the problem into a domain where a specific action is indicated.

Unfortunately, many managers do not adequately distinguish risk from uncertainty. To them risk and uncertainty are both equally bad and they are not clearly separated from one another. They thus confuse these different domains and different modalities of action are also blurred and indistinct.

The author has seen arguments continue for months over the selection of the best algorithms (for sorting, access speed, or efficiency) when simple experiments would have resolved the issues in a few hours.

Similarly, some companies have literally spent years performing AS-IS analysis of systems in preparation of product selection, when such analysis merely delayed a selection process for replacements when any of the identified alternative would have exceeded current system requirements and could have been completed in a few days.

When confusion between uncertainty and risk occurs, relatively simple decision methods become associated with situations that require full, detailed, analysis (like the situation described by Region IV). Sometimes, quick, easy solutions get bogged down in overly complicated decision making processes that are both unnecessary and inappropriate.

Note that extensive analysis is not always necessary even in region IV. There is great evidence that just a little analysis is often enough to lead on to the best course of action (Gladwell 2001) and extensive analysis does not improve the accuracy of our decision making, but rather merely improves our confidence in the choice or selection (Heuer 1999)

When risk and uncertainty are confused with one another, the best course of action can be obfuscated and obscured. When the distinction between risk and uncertainty is not understood, or when decision makers are unwilling to accept any amount of uncertainty or risk, then the process of making decisions becomes are arduous one. Uncertainty and risk often blur and merge, especially when managers are risk-averse or when they cannot accept either in any measure. When this is the case, the decision process bogs down with a wasteful focus on gathering and processing data to make people feel better about their choices, long after the best options may have been identified. Indeed, much of the perception that analysis is needed is driven by non-rational processes (Ariely 2008 and Harrison 1998)

In such cases, where processes of bounded rationality dominate decision making environments, managers basically become incapable of making decisions except when they are very simple ones (Region I - low risk and low uncertainty). When this situation occurs, managers perform analysis activities in regions II, III and IV (when analysis is really only warranted in Region IV where there is both high risk and high uncertainty).

As noted previously, this often leads to the ”paralysis of analysis”, where participants analyze endlessly instead of acting.

Ultimately, managers need to differentiate and distinguish risk from uncertainty in order to recognize when different types of actions are called for.

4 Conclusion

The inability to differentiate between Uncertainty and Risk can lead to unnecessary analysis and preclude appropriate actions. To avoid this, decision makers must understand how Uncertainty and Risk differ from one another and how their different combinations suggest different courses of action.

Managers and decision makers need to learn that it is impossible to eliminate all uncertainty and should focus more on appropriate modalities of action. They need to understand that action is still possible when risk and uncertainty occur alone. Only then will they be able to act confidently and appropriately when either or both are present.

References

Abdellaoui, Mohammed; Hey, John D. (Eds.); Advances in Decision Making Under Risk and Uncertainty; Series: Theory and Decision Library C, Vol. 42 2008, XIV, 242 p. 57 illus. http://www.springer.com/economics/game+theory/book/978-3-540-68436-7 ; Springer-Verlag New York, LLC 2008

Ariely, Dan; Predictably Irrational: The Hidden Forces That Shape Our Decisions; Harpers Collins 2008

Bensoussan, Babette E.; Fleisher, Craig S. Analysis Without Paralysis: 10 Tools to Make Better Strategic Decisions;Pearson Education 2008

Bonini, Charles P., Warren H. Hausman, Harold Bierman. Quantitative Analysis For Management. 9th ed. Boston, MA: Irwin McGraw-Hill, 1997.

Bodily, Samuel E., Robert L. Carraway, Sherwood C. Frey, Jr., Phillip E. Pfeifer. Quantitative Business Analysis: Text and Cases. Boston, MA: Irwin McGraw-Hill, 1998.

Evans, James R. Statistics, Data Analysis, & Decision Modeling. 3rd ed. Upper Saddle River, NJ: Pearson Prentice Hall, 2007

Gladwell, Malcolm, Blink - The Power of Thinking Without Thinking, Litle, Brown and Company, NY, 2005

Harrison E Frank; The Managerial Decision-Making Process; Houghton Miflin Harcout, 1998

Heuer, Richard J; The Psychology of Intelligences; 1999; Center for the Study of Intelligence; CIA; www.odci.gov/csi

Holloway, Charles A. Decision Making Under Uncertainty: Models and Choices. Englewood Cliffs, NJ: Prentice Hall, 1979.

Riabackel, Ari; Managerial Decision Making Under Risk and Uncertainty; AENG International Journal of Computer Science, 32:4, IJCS_32_4_12; Nov 2006 http://www.iaeng.org/IJCS/issues_v32/issue_4/IJCS_32_4_12.pdf

Sewell, Martin; Decision Making Under Risk: A Prescriptive Approach; mvs25@cam.ac.uk; University of Cambridge; Behavioral Finance & Economics Research Symposium – 2009; Chicago 23-25 September 2009 http://www.cs.ucl.ac.uk/staff/M.Sewell/DecisionMakingUnderRisk.pdf