

Return to the “natural” process of decision-making leads to good strategies*

Vladimir A. Masch

RCO Software Limited, Warren, NJ 07059-5453, USA
(e-mail: vladimir@rcosoftware.com)

Abstract. A novel approach to strategic management, *Risk-Constrained Optimization*[®] (RCO), is an ensemble of special models, procedures, and algorithms to generate, evaluate, and help in executing good alternative strategies. RCO is a patented system of planning under uncertainty that searches for the most acceptable compromise between improving results and reducing risk in our decisions. Risk management and scenario planning dominate over optimization. RCO still uses maximization, but only in combination with several protective filters that screen, modify, and scale back the strategies, as necessary. As with any protective equipment, RCO could reduce the need for knowledge about the future.

Keywords: Strategic management – Risk management – Scenario planning – Decision-making under uncertainty – Management science

JEL Classification: D21, D81, C61

1 Introduction

In the current uncertain geopolitical and economic environment, where we cannot reliably know the probabilities of future scenarios, existing theories of strategic business decision-making may lead to serious mistakes. To upgrade the quality of the process, I propose a novel prescriptive approach. It uses an ensemble of special models, procedures, and algorithms to generate, evaluate, and help in executing

* I would like to acknowledge extremely helpful and insightful comments of Professor Mark Perlman (to whom this article owes very, very much), Professor Irma Adelman, Erwin Rezelman, Dr. Victor V. Masch, and an anonymous reviewer. I also greatly appreciate help in editing by David Owens and Professor Irving Rothman.

good alternative strategies. The ensemble must meet very rigorous conditions. It still uses conventional maximization methods, but only in combination with several protective filters that screen, modify, and scale back the strategies, as necessary. As with any protective equipment, such an ensemble could reduce the need for knowledge about the future.

This ensemble is embedded in *Risk-Constrained Optimization*[®] (RCO). RCO is a system of planning under uncertainty that searches for the most acceptable compromise between improving results and reducing risk in our decisions. RCO is made up of three major parts: (a) construction of a large number of scenarios; (b) generation of multiple candidate strategies that are good, flexible, robust, and risk-limited; and (c) identification of a few of the most suitable candidates and subjective final selection of one of them as the strategy to be implemented.

RCO is based on two main ideas. First, the formal methods not reliable enough to find the best alternative strategies are still good enough for an easier task – to screen out the worst. Second, the decision-maker has to be actively involved not only in the final selection of the strategy, but also in the whole process of generation, evaluation, and screening of the alternative candidates.

RCO replicates the “natural” process of decision-making (“multiple everything” – goals, risk types, scenarios, and strategies, with simplified selection of the preferable strategy), not deviating from it by a single unwarranted assumption. More sophisticated utility-maximization decision methods are considered just a special case of the “natural” process. Accordingly, in a more general “natural” framework, these methods are relegated from the top-level conceptual paradigm of the final strategy selection to auxiliary operations, which may, or may not, be used at the earlier stages of constructing and screening the strategies.

RCO is an honest and hard-to-manipulate system that:

- For the first time in more than 50 years legitimates the high-level use of both computers and optimization models in strategic management.
- Creates a protective device that is most reliable, since it screens the strategies by tools unsurpassed in their meticulousness.
- Allows transforming standard tools into realistic customized models, even by executives who have no modeling skills.
- Might provide the tools for implementing the uncertainty-centered decision systems of J. M. Keynes and G. L. S. Shackle.

RCO has been granted a US patent. It could become the “disruptive technology” that business so badly needs to replace the current, greatly deficient “expedient” approach to strategic decision-making – both in its normative theory and in its techniques.

2 The world is changing

In 2000 – before the recent plagues of global terrorism and corporate malfeasance – forty Fortune-200 CEOs were removed. “When 20 percent of the most powerful business leaders in America lose their jobs, something is clearly wrong,” say Bossidy

and Charan (Bossidy and Charan, 2002, p. 14). The authors seem to attribute this phenomenon to the lack of good execution. But perhaps we should look beyond this catchall explanation. Of course, ambition, overconfidence, greed, and sheer ineptitude may have led to many mistakes, but where is a solid and unquestionable methodological counterbalance to the impact of these timeless human failings?

Similarly, “More than 70% of new manufacturing plants in North America . . . close within their first decade of operations. Approximately three-quarters of mergers and acquisitions never pay off” (Lovallo and Kahneman, 2003, p. 58).

Let us look for more detail into one important area of corporate activity – supply chain management. In 1131 supply chain “glitches” reported in the media in 1989–1999, the company’s stock price dropped on average 18.5 percent during the two quarters before and two quarters after the reported disruption. The total destruction of the shareholder value is estimated at between \$151B and \$172B, or on average about \$140 million per glitch (Singhal and Hendricks, 2002, p. 22). (Keep in mind, too, that not all such glitches are reported in the media. Quite the reverse, the information about a company’s supply chain performance is typically kept secret from researchers and the media. There are many large skeletons in corporate closets.) The frequency of these disruptions steadily increased with time – as did the number of the “supply chain optimization systems” in use. A coincidence? As Ian Fleming’s villain Goldfinger said, “Once is happenstance. Twice is coincidence. The third time it’s enemy action” (Fleming, 1959, p. 123).

Of course, eliminating from the forecasts the impact of “planning fallacy” and “delusional optimism” will always be helpful (Lovallo and Kahneman, 2003, p. 58). But these possibilities are limited, and a sound structure cannot be erected on quicksand. One of the main reasons for poor performance of enterprises still is the absence of an adequate foundation, namely, good decision methods. *Business desperately needs valuation and comparison methods for non-trivial, non-associative, strategic decisions – methods that are honest, realistic, constructive, conceptually faultless, and difficult-to-manipulate (with results largely independent of those input data that are easily malleable). At least one of these attributes is missing in every known approach, without exception.*

True, the absence of perfect decision-making methods never prevented the business community from making reasonably good decisions. The existing methods, even if not meeting every high standard, can be considered as “rules of thumb” that provide approximate answers. Through millennia of practical use, only the fittest rules survived. By definition, the survivors should be sufficiently good, and these rules continue to be applicable in relatively stable situations.

But the current nervous geopolitical and economic environment adds a new dimension and urgency to the need for improved methods. On September 11, it became abundantly clear that neither the outcomes nor – especially! – the probabilities of future scenarios can ever be reliably known, and this condition is likely to prevail for the foreseeable future. A decision made when we know (or can subjectively assign) reasonably reliable probabilities of these scenarios is a situation of “insurable risk.” Otherwise, we have “uninsurable uncertainty” (Knight, 1921; Luce and Raiffa, 1957). (I relaxed the requirements here, substituting “reliable” for “objective.”) In the case of “risk,” using expected (weighted average) results

is an acceptable approximation. This has been amply demonstrated in regard to both economic and government decisions. (For instance, Rubin and Weisberg provide a number of high-level application examples of what they call “probabilistic thinking” (Rubin and Weisberg, 2003, pp. xi, 8, 41–43 and 56–57). *Decisions under “uncertainty” follow completely different rules and are incomparably more difficult.*

While our ability to make good decisions appears to have plummeted, risk management and contingency planning become more important than ever, due to our greater vulnerability and the graver consequences of actions.

Less important, but still significant, is the role that uncritical application of “optimization” systems, such as supply chain management, has on decision-making. These systems can be very helpful, but also introduce a substantial risk component of their own, as many corporations are now discovering. (See Sect. 6.) A “successful” implementation of such a system may lead to a supply chain too tight and brittle. To paraphrase an old joke, while a pessimist sees a glass as half empty and an optimist as half full, an optimizer may recommend cutting off the upper half of the glass.

3 Utility maximization: down from decision paradigm to strategy screening

The theories related to decision-making are of three major types (Kleindorfer, 1993, p. 177):

- (a) Descriptive – Study and explain how actual decision-makers perform their activities.
- (b) Normative – Tell how decision-makers should ideally perform these activities, based on abstract models.
- (c) Prescriptive – Help decision-makers to improve their performance, given the complexities and constraints of real life.

The mainstream normative theory of business decision-making is based on the idea of unbounded maximization of expected utility. Both the conceptual soundness and practical applicability of the utility maximization theory are more than doubtful. Among others, we have to take into account the following considerations.

First, any decision-maker has multiple, diverse and conflicting goals and factors of concern (we will uniformly call them “risk types”). Too abstract goals, such as “survival” or “utility”, still have to be defined in terms of a number of more specific risk types. The decision process does not start from an already existing “utility payoff” matrix, with strategies, scenarios, and payoffs provided by somebody else, so that the decision-maker’s task is just to select the “best” strategy by applying to payoffs some aggregation criteria. Payoffs are derivatives. To determine their values, we have to start from the beginning, with finding the values of their precursors, heterogeneous strategy outcomes in each of different risk types, and then convert these values into homogeneous payoffs.

With the conversion “utility functions” both hardly known and subjective, the payoff values clearly are less reliable than the outcomes they stem from. The outcomes are therefore a better basis for strategy comparison. Moreover, the outcomes

allow us to use a simpler, more straightforward, meaningful, and effective strategy evaluation procedure: their values can be compared with some objective or subjective benchmarks and boundaries, thus determining whether a strategy is acceptable, should be rejected, or needs modification.

Since strategy comparison on the basis of outcomes is more reliable than comparison by utility values, we will call the first procedure “strong screening” and the second “weak screening.” Utility becomes less important.

Second, suppose that we successfully calculated the values of “strategy vs. scenario” payoffs. As argued in Section 2, for a foreseeable future we will be in a situation of “uninsurable uncertainty.” Under uncertainty, decision-making is very difficult – it was shown 50 years ago that none of the then known strategy comparison criteria could be deemed “the best,” that is, both theoretically correct and good under any conditions (Luce and Raiffa, 1957, pp. 278–306; see also Section 5(viii)). We have to evaluate strategies by several criteria. Maximization *per se* becomes impossible.

Third, the decision processes in which the concept of homogeneous payoffs and utility maximization can be applied are only a special case of much more general “natural” process of decision-making, used by the animal kingdom for hundreds of millions years, and later – by mankind.

I see this approach as having only three operations: (a) identification or development of decisions or strategies; (b) screening out any strategy that has – under some possible future conditions – a clearly unacceptable outcome in any risk type; (c) subjective final selection of one strategy out of the remaining subset of candidates. None of these operations requires any formal or sophisticated methods. The survivability and universal applicability of the “natural” process are wholly explained by the simplicity of these operations, which require an absolute minimum of intellect.

The “natural” approach does not therefore rely on any unwarranted assumptions and incorporates “multiple everything” – strategies, scenarios, and different risk types that include goals and may be non-formal. Outcomes for heterogeneous risk types do not have to be converted into homogenous payoffs, and these payoffs (if they are calculated) do not have to be maximized.

The “natural” approach is common in non-business decision-making. (Any decision problem can be presented in this framework. For instance, we can formulate several strategies, such as “to take an umbrella,” “not to take an umbrella,” or “to buy a pocket umbrella,” and two or more scenarios describing the possible weather.) It is abandoned in business because of the increased complexity of constructing both scenarios and strategies and calculating the outcomes.

If, however, the decision problems solvable by the utility maximization methods are just a narrow special case, then – for the more general class of problems – these methods cannot constitute the basis for the normative, top-level conceptual paradigm of selecting a strategy. Their proper place should be in the auxiliary operations of constructing and screening the strategies, while the final selection of the strategy is performed by non-formal methods. We do not lose the advantages of using – whenever possible – elegant and efficient maximization models and computers; we simply apply them in another, proper, better place.

Fourth, I think that a recently proven result casts additional doubt on the very soundness of the utility maximization criterion. Radner and Shepp have devised a beautiful and simple axiomatic model to study optimal decision-making for a firm with uncertain earnings that maximizes its profits in a certain natural sense – by its choice of both corporate policy and the timing of withdrawal of its profits. They conclude that, if such a firm follows its optimal profit maximization policy, it will reach bankruptcy at some finite time, with probability one (Radner and Shepp, 1996). Shepp (private communication) has already developed a simple formula to compute the expected time to bankruptcy under the optimal policy within this model.

In distinction from these two authors, who pursue profit maximization and therefore do not see their results as a negative, I think that the proper goal of a business is long-term survival (where profit is not a goal, but a constraint); profit maximization is shown to be inconsistent with it. Of course, for a business, profit is not a utility, but it still is a major component of it.

Finally, the current normative model, which proposes pure offense (maximization) without defense (risk management), is wrong. It would not be accepted in sport, by the military, even by gamblers. Why should it be acceptable for business?

4 The new approach

But let us forget about the legitimacy of utility maximization as a leading theoretical principle. Irrespective of whether one believes in it or not, the Holy Grail of decision-making would be to find a prescriptive ensemble of methods to embed the conventional maximization models that would make these models safer to use, more productive and more reliable. *As with any protective equipment, such an ensemble could reduce the user's need for knowledge about the future; hence its critical importance.*

Paradoxically, that Holy Grail can be found only after a full and undisguised return to non-normative subjectivity, which presumably should lead to chaos, but actually would free us to “create order out of chaos.” That is, we can safely use powerful maximization models and methods to develop far-reaching “optimal” but risky candidate strategies, if, getting rid of all illusions, we use them only as one part of the process. In other parts, we carefully screen, modify and scale back these candidates, to protect the decision-makers, under any future conditions, from unacceptable outcomes in any risk type, and thus derive good, robust, and limited-risk strategies. Risk management, imposed in accordance with the risk attitude of the executive, should dominate over the maximization of rewards. That is, business decision-making should, like medicine, follow the principle “First, do no harm.” This principle is not so bad from the business point of view, either. For instance, it closely matches “the first ground rule” of a prominent British businessman: “Assess risk rather than reward. Don’t worry about the upside: that will take care of itself.” (Voyle, 2002).

Using a sports analogy, maximization models become our “offensive line,” which we can use now most aggressively because it is complemented by several protective filters, our “defensive lines.” Some of the filters should be general, rather

than model-specific, to allow us to deal with flaws not only in the information available, but also in the very maximization models used.

Also, we should not limit our choice of filters to a single field of science. A radical change in the mode of business decision-making may become a “disruptive technology” that will have no less impact than a major industrial innovation. And “. . . when an idea moves from an invention to an innovation, diverse ‘component technologies’ come together. Emerging from isolated developments in separate fields of research, these components gradually form an ‘ensemble of technologies’ that are critical to each others’ success.” (Senge, 1990, p. 6).

Computers and all formal operations can now be used not to dictate, but to perform two interconnected tasks:

- Presenting the decision-maker with a small screened set of good alternative candidate strategies, none of which can lead to a disaster;
- Informing the decision-maker fully about the behavior of each candidate strategy under the whole range of scenarios in regard to each risk type, as well as about the contingency plans that might arise from the strategy.

Decision-making becomes an iterative “man-computer” process, where the capabilities of both participants are used properly and efficiently. Computers finally find their place as legitimate analytical tools in high-level strategic decisions, while the human decision-maker directly controls the construction, screening, and modification of the candidate strategies and, in the end, makes a subjective but fully informed and reasonably safe final strategy selection. (The current use of computers in strategic management does not rise above the data processing precursors to decision-making, even if sophisticated and glorified precursors, such as simulation.)

The proposed approach also changes how we find out the expected personal response of an executive to a problem. Currently it is like a “two-way translation”: first, an analyst attempts to simulate the executive’s risk attitude and subjective preferences, and, second, we derive the expected response from that simulation. Since, under the proposed approach, the decision-maker becomes an active participant in the process, we no longer need such a roundabout, costly, and error-prone approach, where we never know, for instance, if the revealed risk attitude refers to his spending his own or someone’s else money. Instead, we can discover the decision-maker’s response by simply asking him “Is this strategy outcome acceptable?”

All techniques outlined in this section are means to achieve one end, which is the long-term sustainability of the business. To summarize, I propose an approach to business decision-making that must:

- Strive for a reasonably good strategy, with emphasis on long-term sustainability and protection from risk in all heterogeneous risk types, in accordance with the risk attitude of the decision-maker. Provide the ability to survive and to respond, whatever the future brings about.
- Discard all claims to the process following a normative ideal and fully accept the unavoidable primacy of subjective judgment. Relegate the formal methods, models and formulas to a less demanding role of either providing negative recommendations (that is, filtering out bad or risky strategies) or being used to create purely analytical inputs to decision-making.

- Discard all unwarranted assumptions about the process of decision-making. Make it consistent with proven concepts and results of psychology, behavioral economics, the “satisficing” approach, and a comprehensive “natural” approach (as described above). Greatly diminish reliance on probabilities of future scenarios.
- Bring the decision-maker directly into the midst of the strategy generation, evaluation, and screening process, and thus eliminate the need for auxiliary models that allegedly simulate his risk attitude and subjective preferences. Create realistic models exactly tailored to the situation, to the attitudes and preferences of the decision-maker, and to the intended trade-offs.
- For each candidate strategy, develop contingency plans over the whole range of forecast scenarios, and evaluate each candidate strategy on the totality of outcomes over all these plans.
- Screen multiple candidate strategies through the use of several filters to provide to the decision-maker a small subset of good and reasonably safe finalist candidates, as well as descriptions of their behavior under different conditions.
- Make strategic decision-making an iterative and collaborative “man-computer” process, accepting the primacy of the human decision-maker, where the capabilities of both participants are used properly and efficiently.
- Be able to overcome major deficiencies in both information about the future and in the models used, incorporating for that purpose an ensemble of “component technologies” that come from diverse fields of science.

The proposed approach should unify the process of planning and decision-making – with the planners and decision-makers, on the one hand, and the result of that process (a decision, a plan, or a strategy) on the other hand. This should stop the disputes between the three schools of business planning – the rationalist, the evolutionary, and the processual. (M. E. Porter is a rationalist, H. Mintzberg is an evolutionist, and General D. D. Eisenhower, who allegedly has said “Plans are nothing. Planning is everything,” is a processualist.)

I define “*business wisdom*” as *the capability to generate, to evaluate, and to execute good alternative strategies*. “Business acumen” cannot replace “business wisdom.” Perhaps one possible answer to what was wrong with the forty CEO’s removed from their positions (see Sect. 2) was that, too often, they used their acumen instead of their wisdom and “expedient” decision-making instead of “reflective” (Jehn and Weigelt, 2001). The evolutionary school of planning considers this conduct the norm; I disagree. It is better to design a limited-risk strategy earlier and to have good contingency plans ready, than to fight fires haphazardly after they erupt.

5 Risk-constrained optimization[®]

(i) General comments

A system that meets all rigorous requirements of Section 4 is Risk-Constrained Optimization[®], or RCO. RCO is a system of planning under uncertainty that searches for the most acceptable (not “the best”!) compromise between improving

results and reducing risk in our decisions. RCO has been granted a United States patent (US Patent, 1999).

“Risk-constrained optimization” is no more “optimization” than “bounded rationality” is economic “rationality”; the qualifiers are weightier here than the concepts being qualified (Simon, 1982). Protection from risk dominates in RCO over maximization. But RCO can include such maximization tools as mathematical programming models, because their output is filtered through five diverse layers of risk protection. The main version of RCO includes the following stages:

- Constructing a large number of scenarios of the future and performing the initial “What if” analysis.
- Formulating a multiscenario mathematical programming model and using the model solution as a starting point for an iterative process of generating a number of alternative candidate strategies.
- Recording the scenario outcomes for any candidate strategy in a three-dimensional “outcome matrix”: “scenarios vs. strategies vs. risk types.”
- Obtaining a number of alternative candidate strategies by re-running the model, in an iterative process, each time with a different set of “risk-limiting constraints” imposed on any unacceptable outcomes.
- Converting the multidimensional “outcome matrix” with heterogeneous outcomes into a two-dimensional “payoff matrix” with homogeneous payoffs: “scenarios vs. strategies.”
- Screening out the worst alternative candidate strategies by joining several decision criteria in a framework of “strategic frontier,” leaving for final consideration a small subset of the best and safest candidate strategies.
- Making from that subset a final, subjective, but well-informed and reasonably safe, selection of the strategy to be implemented.
- Deriving for the selected strategy a set of contingency plans over the whole range of scenarios.

As noted in Section 3, the five types of filters that screen the strategies are of different quality. The three first filters (the multiscenario model, the outcome matrix, and the “risk-limiting constraints”) deal directly with the non-converted outcomes of heterogeneous risk types and are therefore “strong” filters. The remaining two filters (using several decision criteria and applying “strategic frontier”) deal with utilities and are “weak” filters.

(ii) Scenarios and the “what if” analysis

The process starts with the generation of multiple scenarios of the future. RCO scenarios are different probable combinations of (a) single values of “predictable” parameters, and (b) alternative values of those model parameters whose values we cannot reliably predict. These combinations must be internally consistent. (Correlation of parameters is taken into account by (RiskMetrics, 1996), a well-known methodology.) There may be a very large number of scenarios. This is not a big obstacle, because RCO has special computational methods intended for problems

of that type and size, which may involve clustering (and de-clustering) scenarios, as described in (iv).

At this stage, we can discover whether the extended RCO process is needed or not. We construct a prototype deterministic (single-scenario) optimization model and then diversify it by scenarios, putting in each “what if” model the parameter values that are specific to the corresponding scenario. Each “what if” problem is considered here as independent, that is, not connected with other “what if” problems. If the solutions of these problems lead to a similar strategy over the whole range of scenarios, and this strategy is completely satisfactory in all risk types, we may finish the process here. However, this is unlikely.

(iii) The multiscenario model

In the beginning of the strategy generation part, RCO constructs a multiscenario stochastic programming model, which, in contrast to stage (ii), considers the whole range of scenarios jointly. The multiscenario models have a great advantage over the type (ii) single-scenario models: by definition, the latter are not able to find any “interscenario compromise” solutions, which are correct answers for most problems.

For example, suppose that we have to build only one factory to manufacture a product that would be in demand solely at one coast of the USA, either Atlantic or Pacific, and that the cost of transportation of this product is very high. Suppose also that we have to build the factory before we know where the demand would be. Under these conditions, we should locate the factory somewhere in the middle between the coasts. Even the simplest two-scenario model, with demand on both coasts, easily finds this solution. But no single-scenario optimization model would ever be able to find it, since each of them is tied to just one of the coasts.

The RCO multiscenario model is of a new special type. (Stochastic models are used in OR/MS for planning under uncertainty, but their currently existing types are not sufficient for RCO purposes.) The main distinction of the RCO model is that it recognizes that not all decisions are created equal. RCO therefore classifies all decision variables of the model into “strategic” and “operational” variables. “Strategic variables” are a small subset of variables that correspond to the most important decisions. These include, but are not limited to, immediate and irrevocable decisions. In the model solution, the value of any strategic variable remains unchanged, or stays within specified bounds, under each scenario. A “strategy” is defined by the set of solution values of all strategic variables of the model. For instance, if the model has only one strategic variable, which is the capacity of a factory to be built, Strategy A might correspond to zero capacity (under all scenarios), Strategy B to capacity from 100 to 105 units, and strategy C to capacity from 200 to 205 units.

Regarding the second distinction from the current stochastic programming, which is interested only in the total weighted value of the objective function of its model over the whole range of scenarios, RCO calculates the “bookkeeping” scenario components of this function, as well as the scenario-specific outcome values of each risk type, such as market share or environmental damage. RCO would need all these values at stage (v).

In a multiscenario model, scenario inputs and results have to be weighted. The scenario weights (which are unreliable and are overridden later) are derived here by combining the available objective and subjective probabilities of future events.

It must be emphasized that RCO works on the basis of incomplete, unreliable, fuzzy, non-quantifiable, and contradictory information. We should not expect accuracy from forecasts and other input data used in the models. Most importantly, RCO should not treat the input data as reliable information. The immediate consequence of that premise is that the methodology should alleviate, as much as possible, the impact of the poor quality of the input data on the quality of solutions. In other words, the methodology should be, to some degree, self-correcting. (A second consequence is that both the computational complexity of the models and the speed and accuracy of algorithms used to solve them should be commensurate with the quality of the data.)

Not only RCO, but all multiscenario models are, to some degree, self-correcting. If some scenarios are faulty and some are correct, then the model solution, by virtue of being a compromise between solutions for individual scenarios, should reflect only partially the impact of faulty data. Also, if a scenario is too optimistic, the outcomes for this scenario will improve more or less in parallel for all candidate strategies, and vice versa. Since what counts is the selected strategy and not the absolute values of outcomes, this shift affects neither the relative attractiveness of the strategies nor the final decision. The multiscenario model thus becomes the first filter of RCO.

During strategy generation, models produce such results as, for example, where and when to build the new factories. But at this stage it is more important that they also provide valuable information (profit, market share, environmental hazards, etc.) about the behavior of a number of good candidate strategies, both in the overall solution and under different future conditions. Looking at all results and using the risk-limiting constraints, the decision-maker is able at stage (v) to change the strategies to suit his preferences and risk attitude, as applied to different risk types.

It should also be emphasized that the multiscenario models provide the “post-contingency plan” outcomes, that is, the strategy’s reactions to various future conditions. The results of a strategy are defined by the totality of outcomes of its contingency plans over all scenarios.

Finding the extremal solution of this model is the end of the stochastic programming process as it is practiced now, but it is just the beginning of the iterative RCO process.

(iv) The “outcome matrix”

The scenario values of the objective function and the outcomes in all risk types are recorded in a multidimensional “outcome matrix,” which is the second filter of RCO. This matrix presents, perhaps in the most condensed format possible, all valuable information which the decision-maker might want to know about the behavior of each candidate strategy under the whole range of scenarios, in regard to all risk types. The matrix is very important, since it allows a “bird’s eye view” of everything and immediate reaction. It is the best way to extract meaning from an

ocean of information. Without it, imposing the “risk-limiting constraints” would be impossible.

Of course, it also would not be reasonable to analyze this matrix for a very large number of scenarios. Therefore, wherever necessary, RCO clusters the scenarios in homogeneous groups. But, to avoid possible clustering errors and missing the risks under bad scenarios, RCO uses de-clustering, which probably would be needed, however, only for a few groups.

At this point of the process, the matrix usually has three dimensions: “scenarios vs. strategies vs. risk types.”

(v) *“Risk-limiting constraints”*

If, after any iteration, the decision-maker is not satisfied with any outcome for any risk type under any scenario or group of scenarios, he adds to the multiscenario model some “risk-limiting constraints,” imposed directly on the poor outcomes. The constraints should improve these outcomes. The more additional constraints the model has, the more its solution differs from the initial solution of the process, that is, from the solution of the multiscenario stochastic model.

Since the constraints are imposed not on the “structural” variables of the model, but rather on the “bookkeeping” variables that reflect only the model results or outcomes, the decision-maker needs no modeling skills to impose these constraints; neither must he understand the internal structure of the model. It is quite sufficient to know the risk types and scenarios.

Suppose that the profit obtained by the latest strategy under a recession scenario (say, Scenario 5) equals \$1,200,000. We would like to find a strategy that increases that profit to \$2,000,000. So we just fire off a command: “profit[5] \geq 2000000.” Now the model will implicitly split all possible strategies into two categories: the “good” strategies that satisfy that condition, and the “bad” ones that do not. The constraint excludes from consideration all “bad” strategies, but, remarkably, does not exclude a single “good” strategy. The previous “optimal” solution certainly belongs among the “bad” strategies and is excluded.

The model is re-run and a new solution is found with the poor outcomes improved as much as possible, reaching the requested desirable levels if they are feasible. Of course, such partial improvements demand sacrifices in other parts of the solution, so the solution as a whole worsens – its total value of the objective function is less than that for the previous solution. It now is no better than that value for the previous “second best” solution. This “Dutch auction” process for a single candidate strategy is complete when either the decision-maker is satisfied or further improvement becomes impossible. After the last constraint for this candidate is imposed, the new region of feasible solutions is some truncated, “safe” inner portion of the initial region of feasible solutions.

Within a trade-off, from one constraint to another, the definition of the “good” and “bad” strategies constantly changes. If there exist no strategies with “profit[5] \geq 2000000,” we may grope for a feasible answer, asking the model if there exist any strategies with “profit[5] \geq 1800000” and what sacrifices are needed in various operations or in other scenarios to reach that level. On the other hand, if

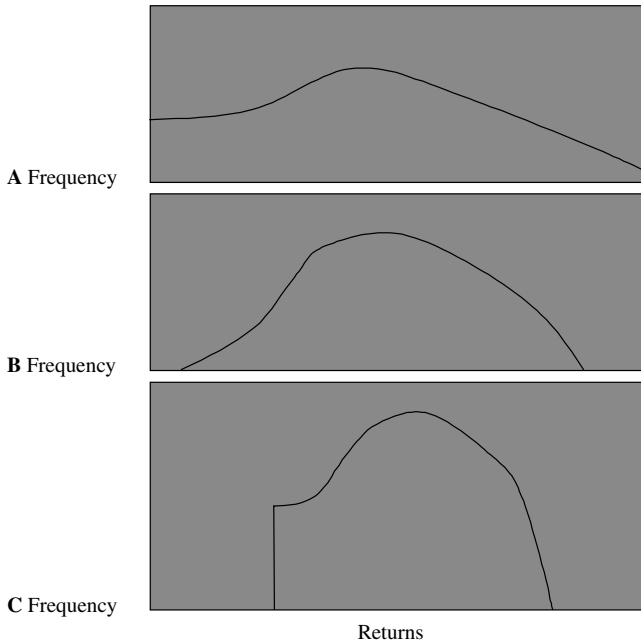


Fig. 1A–C. Distribution of outcomes over the range of scenarios. **A** Optimal solutions of single-scenario models. **B** Compromise strategy, optimal for a multiscenario model. **C** Strategy B after “risk-limiting constraints” are added

\$2,000,000 is reachable, we may try for \$2,100,000 – or, say, for a higher market share. Each branching of a set of constraints leads to a different trade-off.

The risk-limiting constraints imposed upon the multiscenario model reflect both the decision-maker’s propensity for risk and his attitude toward various risk types and optimality criteria. The strategy generation process can be repeated many times, branching out at different points, on different risk types, with different types of constraints, or with different bound values, and resulting in one or more “trees” of candidate strategies with different trade-offs.

Figure 1 illustrates how the risk-limiting constraints work in reducing risk. It shows the effect of RCO on the distribution of outcomes over the whole range of scenarios. “A” summarizes the solutions of single-scenario “What if” linear programming models over that range. “B” shows how the situation changes with the introduction of a multiscenario model, which generates a single compromise strategy. “C” depicts the same strategy after risk-limiting constraints are added.

In “A,” single-scenario models have no mutual impact, so the good scenarios lead to the highest returns and the bad ones to the lowest. As a compromise, “B” has more modest results for good scenarios and is also less risky. Finally, the risk-limiting constraints in “C” cut off all elements that lead to poor outcomes under bad scenarios, so the distribution of returns is truncated on the left. The distribution in “C” is greatly compressed, and, extremely important, especially for Wall Street, the volatility of returns is decreased.

Systematic use of risk-limiting constraints, especially the constraints based on scenario-specific outcomes, is a novel OR/MS technique introduced in RCO. The constraints are a very important third filter of RCO and perform several functions. They “strongly” filter out all unacceptably risky decisions, while not excluding (as shown above) any good strategy, and perform the function of risk management – full or partial prevention of undesirable outcomes at any moment within the planning period.

(vi) The analysis of “shifting probabilities”

One of the major functions of these constraints is to override the initial scenario weights. As indicated above, the scenario weights of the initial multiscenario model include subjective “guesstimates” of the likelihood of future events. In the present environment of high uncertainty, these are more guesses than estimates and are completely unreliable. Clearly, these have to be modified. The RCO process repeatedly overrides and changes the implied scenario weights as it imposes each new risk-limiting constraint.

In that process, it turns out that shifts in the assumed probabilities bring about different optimal solutions and outcomes. Looking at the likelihood and desirability of these outcomes helps us, retroactively, to clarify our perceptions of what our level of “optimism” is in the specific situation we face. Such estimates would be extremely helpful later, at the stage of applying the Decision Science criteria; see (viii)–(ix) below and an example in Section 9. This novel type of analysis is intertwined with the imposition of the risk-limiting constraints and effectively contributes to their filtering function.

At this stage, the outcome matrix adds a fourth dimension to become a “scenarios vs. strategies vs. risk types vs. probability ranges” matrix.

(vii) Calculation of “utility”

At the next stage (viii), we compare strategies by applying Decision Science criteria to prune the list of candidate strategies. All these criteria, however, are designed for a two-dimensional “payoff matrix” with homogeneous payoffs, in contrast to an “outcome matrix” with three or more dimensions and heterogeneous outcomes. In other words, if we want additional “weak” screening, we still have to determine the “utility” of such different and complex factors as market share, consumer attitude, environmental hazard, or a company’s workforce stability. Depending on the risk type, RCO has several novel methods of conversion. The simplest (and universal) method is to use two or more conversion coefficients for each risk type, thus increasing the number of scenarios we have to consider. (For instance, suppose that we derived a utility function that equates one percent of market share to 500,000 of utility units. Instead, we can use three values of 450,000, 500,000 and 550,000 units.) At this stage, RCO can easily handle any number of scenarios.

(viii) Decision criteria

Each candidate strategy is defined now by a set of “payoffs” – utility values under the whole range of scenarios. To select a strategy, we need to compare the strategy values of some range aggregate of these payoffs. There are four major possible range aggregates for each strategy: (a) the weighted average payoff and (b) the non-weighted average payoff over all scenarios, (c) the maximum (“best case”) payoff under a single scenario, and (d) the minimum (“worst case”) payoff under a single scenario. There are four similar “regret” measures as well. (“Regret” is a derivative of payoff that measures the opportunity lost. For each scenario, the regret for Strategy A is calculated as the difference between the best payoff under that scenario and the payoff of Strategy A. Thus, if the best payoff under scenario 1 – say, for Strategy B – equals 100, while the payoff for A is 10, then the regret for A under scenario 1 is 90: if we choose Strategy A and scenario 1, indeed, occurs, we would have regret of 90 units that we have not selected Strategy B instead of Strategy A.)

RCO pairs these eight “single” criteria into six “synthetic” criteria, three for payoffs and three for regrets. A “synthetic” criterion means a weighted average between a more optimistic and more pessimistic outlook, the latter being represented by the “worst case” payoff or regret. The weights used in these averages reflect our subjective evaluation of the chances of “good vs. bad” outcomes, which approximates the risk attitude. They also reflect our trust in the reliability of the input data. The payoff criteria are: “index of pessimism-optimism” (“PO index”), that considers an average between the “best case” and the “worst case” payoffs; “partial ignorance index” (“PI index”), for an average between the weighted average payoffs and the “worst case” payoffs; and the “insufficient reason index” (“IR index”), for an average between the non-weighted average payoffs and the “worst case” payoffs. The regret criteria are similar, with the regret values instead of payoffs. (These criteria can be further expanded.) “Single” criteria thus become merely the narrow special cases of more comprehensive and meaningful “synthetic” criteria. The “PO index” criterion of L. Hurwicz was developed earlier (Luce and Raiffa, 1957, p. 282; Hurwicz, 1951); other “synthetic” criteria had been introduced by RCO. Two of the “synthetic” criteria are demonstrated in the example in Section 9.

As noted in Section 2, the “single” weighted average criteria can be used for decisions under risk, when scenario probabilities are considered reliable. Under uncertainty, no single criterion is “the best.” This conclusion has been only strengthened by our introduction of multiple “synthetic” criteria. In other words, any problem may have several “best” solutions, based on different criteria. Therefore RCO uses all six “synthetic” criteria – jointly, but in turn; “single” criteria are not used. (The “synthetic” criteria can also be jointly used under risk, instead of the weighted averages, providing a more versatile and reliable comparison.)

Six “synthetic” criteria provide for a “weak”, but still very important fourth filter of RCO. In contradistinction to the risk-limiting constraints (the third RCO filter), these methods are not connected with the “optimization” model used to generate the candidate strategies. Therefore, they help compensate, to some extent, for possible flaws in the model.

The crucial advantage of the “synthetic” criteria is that four out of six do not need scenario probabilities at all, while the other two are only weakly connected with them. This frees RCO from almost any reliance on probabilities, which are the weakest part of the input data. Instead, RCO focuses on *possibilities*.

This also makes RCO an honest system: manipulation of probabilities is the favorite technique for distorting the comparison of the strategies. It is much easier to change subjective probability than, say, the cost of a contingency measure. Choosing a wrong range of scenarios is a variation of the same manipulative technique: it is equivalent to assigning zero probabilities to undesirable scenarios. That is what Enron used to do – when making a strategic decision, it thoroughly analyzed up to 1,000 scenarios (DeLoach, 2000, p. 194). All scenarios were taken, however, from an overly optimistic range.

An additional benefit of “synthetic” criteria is that, like many of the “mass customization” technologies used in industry, they offer the possibility of delaying the most difficult (or, in our case, the most controversial) operation until the very end of the process, or until such a time when its introduction will inflict the least harm. “The most controversial” operation here would be the use of subjective estimates of the decision-maker’s optimism.

(ix) The “strategic frontier”

Moreover, RCO uses the “synthetic criteria” in such a way that they become much more than just a better technique for comparing strategies. They are embedded into a framework of a new “strategic frontier,” which is the fifth filter of RCO, also not connected with the “optimization” model. This “strategic frontier” could not be applied with any “single” criteria.

The frontier incorporates valuable information about the relative merits and faults of any strategy:

- The composition of the subset of strategies that form the frontier.
- The width of the interval supporting each frontier strategy.
- The order of the frontier strategies on the optimism-pessimism spectrum.
- The difference between the frontier strategy and each other strategy, which shows the possible impairment of results in choosing a non-frontier strategy.

The “strategic frontier” allows us to apply subjective estimates in a more prudent, convenient, and less demanding way (that is, the decision-maker does not need to specify in advance his degree of optimism). Specifically, the frontier replaces hard-to-estimate “point” indices by index ranges. For instance, the current user of the “PO index” criterion may ask the question: “Which strategy, B or C, is better if our estimated value of the index equals 0.8?” This means that we compare the strategies at precisely 0.8 probability of the “good” outcome and 0.2 of the “bad” outcome. Instead, when we use the strategic frontier, it is sufficient to say that Strategy C is preferable if the probability of a “good” outcome is no more than, say, 0.889, and Strategy B otherwise. Again, the strategic frontiers for two criteria are demonstrated in Section 9.

The “synthetic” criteria and strategic frontiers do not select the strategy to be implemented. They just shrink the list of the “finalist candidates” to a few best and safest strategies, leaving the selection to the decision-maker. Similar to the risk-limiting constraints, they do not find the best – they eliminate the worst, which is easier.

(x) The final selection and the contingency plans

The decision-maker eventually selects one strategy, based on the most acceptable combination of the outcomes for all risk types under all scenarios, the value of the objective function being one of these types. The model solution for this strategy is then used to derive more detailed and specific information, including contingency plans for each scenario. We already have these plans as the scenario components of that solution, so we have only to extract and fully utilize these results. As mentioned before, the totality of outcomes for the contingency plans over the whole range of scenarios is the basis for evaluating each strategy. *Having good contingency plans ready in advance of any thinkable crisis is the RCO’s contribution to executing the selected strategy.*

6 Customized models

From its very inception in 1947, linear programming was described as an activity aimed at the “optimum,” or “best,” allocation of limited resources. “Optimization” has a strong, positive connotation

Indeed, linear programming algorithms run at a speed faster than lightning, so solvable models can be of truly enormous size. They can include millions of constraints and scores of millions of variables (compare that, say, with statistical models, where the possible number of variables rarely exceeds several dozen). Mathematical programming models are the only way to integrate the planning needs of a large company. “To integrate” means here to interrelate, connect, balance, synchronize, and jointly analyze different factors, operations, and territorial or functional parts of a business. These models and algorithms are also unbelievably good in ferreting out the slightest possibility of moving the solution of the problem in the desired direction. In short, they are extremely powerful – the very zenith of computational efficiency.

And yet, like computers, they never have been applied successfully in high-level strategic management. Leading journals, such as *Harvard Business Review*, stopped publishing articles on optimization more than twenty years ago, citing a lack of interest among business leaders.

The main reason for that failure is obvious: the more powerful a tool is, the more dangerous it may become in unskilled hands or when wrongly applied. In fact, the “optimization” models do not find the *optimal* solution of the real problem that faces the decision-maker. Instead, they find the *extremal* feasible solution of the mathematical programming model that supposedly describes the real life problem, so that “optimization” (which should be good) becomes a euphemism for “extremization” (which is likely to be bad). In non-trivial problems, those suppositions of

model realism are seldom if ever true, and the models can be very far removed from reality. Since the solutions of such models are very unstable and easily switch from one extreme to another, even a minuscule deviation from reality may lead to veritable disaster. These models thus carry a substantial risk component of their own. In the present mode of application, they are safe to apply only if their “cages of constraints” are so restrictive that the solution cannot deviate much from the status quo in any direction, including the wrong direction. Almost by definition, strategic management requires a deviation from the status quo, so using such models here is impossible. As one of the founders of the discipline said, it “is capable only to rearrange the chairs on the deck of the Titanic.”

Of course, no model can fully reflect reality. But the model must reflect those important features of reality that would move the decision in a proper direction. What is troubling here is that, even though an analyst knows what factors and considerations he incorporated in the model, he has no idea at all how important are those factors and considerations that he has omitted. Obviously, he considers them to be insignificant, but they still may be of paramount importance. These factors and considerations can be separated into two major groups:

- a. “Enterprise-specific” – not incorporated because they are hard to formulate and quantify. Also, the decision-maker might know some of them from his practical experience, but not the modeler. Examples: long-term considerations in a short-term model; model components that express “the system of beliefs” of the company’s management (Donaldson and Lorsch, 1983); and stability of the physical processes, employment, and other components of the enterprise’s activities. The threat of terrorist acts – with their location, scope, outcomes, and especially probabilities, unknown – has lately become perhaps the most prominent element of this group.
- b. “Situation-specific” – probably are not known in advance at all and can be found only in the process of applying the model. They include the personal attitudes and preferences of the decision-maker and should be incorporated as we learn about them during the process. For instance, we might see what level of losses, in which operations and situations, is unacceptable to the decision-maker.

We probably will not be able to model well most factors and considerations of the first group, except “the system of beliefs.” It is more important, therefore, to incorporate the second group, which is done by “risk-limiting constraints.” They are our best possible approximation to defining the acceptability of a strategy and to rendering the model realistic and relevant.

The combined impact of all these model deficiencies should be so great that we perhaps may state that, in any non-trivial problem, a standard model, or even a model customized by an analyst (without a direct participation of the decision-maker), is not sufficiently reliable to provide good recommendations.

Any non-trivial decision model should be “tailored” – adapted to the situation, to the decision-maker, even to each specific trade-off we are trying to achieve. The decision-maker should be directly and actively involved in this adaptation. To assure efficiency of the tailored model, it should be constructed iteratively, in a

Dutch auction manner, so that we get the best answer for the specific tradeoff that remains after all constraints are met. RCO performs all these functions precisely.

7 The “ensemble of technologies”

RCO merges technologies that belong to six fields: (1) Operations Research/ Management Science (OR/MS), (2) Scenario Planning, (3) Decision Science, (4) Risk Management, (5) Utility Theory and (6) Portfolio Theory. Scenarios come from OR/MS, Scenario Planning and Decision Science. Improved multiscenario models, as well as novel “risk-limiting constraints” and computational methods, belong to OR/MS. The “analysis of shifting probabilities” belongs to both OR/MS and Decision Science. The novel multidimensional outcome matrices and “synthetic” decision criteria belong to Decision Science. The conversion of a multidimensional outcome matrix into a two-dimensional payoff matrix is a contribution to Utility Theory. The novel strategic frontiers contribute to Decision Science and Portfolio Theory. RCO is closely interwoven with Risk Management. Each of these disciplines required substantial innovation and overhaul. All components above are greatly improved and often are completely new. *The power of RCO comes, however, not from the individual components, but from their combined impact within the process.* In the quoted above words of P. Senge, they form an inseparable ensemble and “. . . are critical to each others’ success” (Senge, 1990, p. 6).

In three of these fields – OR/MS, Scenario Planning, and Risk Management – RCO turns the current approach inside out. The most dramatic change is in *OR/MS*, where RCO uses mathematical programming and other models primarily to eliminate the worst, rather than select the best, decision or strategy, since it is easier to decide what we *do not want* to do, rather than what we *want* to do. In a Dutch auction manner, RCO gradually tightens its set of constraints until all concerns are met for all risk types, as far as possible, to obtain a model that leads to one of the feasible trade-offs. (A similar procedure is also applied in filtering the strategies by the “synthetic” criteria and strategic frontiers.)

This 180-degree change in *modus operandi* kills six plump birds with one stone. The resulting system:

- Creates a big part of a protective gear that reduces the need for knowledge about the future.
- For the first time, legitimates the use of computers in high-level strategic management, which was impossible under the old mode.
- For the first time, legitimates the use of the optimization (mathematical programming) models in high-level strategic management.
- Creates a protective device that is most reliable, since it screens the strategies by tools unsurpassed in their meticulousness, such as the optimization models, “synthetic” decision criteria, and strategic frontiers.
- Dovetails perfectly with the requirements of the “natural” process of decision-making, where the formal methods are needed for developing and preliminary screening of the strategies, rather than for selecting “the best.”
- Allows developing realistic customized models as described in Section 6.

The second field is *Scenario Planning*. In most of the current practice, it ends with just a few “consensus scenarios” or, at best, with a general outline of a strategy. In some cases, these results may be sufficient. Otherwise, the effort deteriorates from “scenario planning” to “scenario analysis.” Anyway, the results usually fail to compensate for many months of extensive effort by many managers, plus costly fees. It is no wonder that most of the conventional scenario planning industry is now in deep trouble (London, 2003).

A strategy can be properly evaluated only on the totality of outcomes under the whole range of scenarios – not of its direct outcomes, but of outcomes of its contingency plans. The devil is in the details; this is particularly true for risk management. Looking at a large number of scenarios is, as a rule, necessary not to miss a potential danger, especially in non-linear systems.

In contradistinction from the current Scenario Planning, RCO has to provide a large number of scenarios, a detailed strategy to pursue, and contingency plans to match. It also must provide comprehensive, specific, numerical results, which can go as deep and as broad as desired. It must perform rigorous economic, financial, and mathematical analyses of both assumptions and results. It should give the decision-makers substantive, cogent, detailed, and defensible analyses of the proposed strategy and contingency plans. To achieve all that, the RCO scenarios must be developed not as descriptive stories, but rather as combinations of different values of the model parameters; see Section 5(ii). Also in contradistinction from the current Scenario Planning, constructing scenarios most usually takes only a small fraction of the total effort and time.

RCO also demands radical change of direction in the third field, *Risk Management*. According to a recent survey, 58 percent of risk events that caused large stock drops in 1995–1998 (that is, before the current geopolitical and corporate malfeasance turmoil) have strategic, and six percent, supply chain origin (Miccolis, 2001, p. xxix). (In financial terms, rather than in the numbers of events, the percentage of strategic risks is even higher.) The types of dangers associated with the present geopolitical situation have further increased the relative importance of strategic risks. This fits the RCO approach, which tries to eliminate or mitigate risks early, by developing a flexible and robust strategy. *In this “egg or chicken” situation, risk management considerations should lead to a strategy, not vice versa.* Instead, both its current practices and publications are completely focused on non-strategic, small-scale improvements – many of which, moreover, become unrealistic now because of the tightening of the available insurance options. These improvements amount to “patching” the already existing strategy by adding to it some elements of possible contingency plans. Creating a hierarchy of “risk owners” able to deal merely with “manageable” risks, as well as centering attention on such backward-looking techniques as “risk mapping,” only perpetuates that harmful trend.

8 The origins

Two great economists focused on the crucial, central role of uncertainty – J.M. Keynes and G.L.S. Shackle. Professor Shackle followed in the footsteps of Keynes and further expanded his approach to decision-making under uncertainty. Af-

ter reading the draft of this article, Professor Mark Perlman, one of the closest comrades-in-arms of Shackle, kindly suggested that RCO is “an implementation of Shackle’s uncertainty formulation.” I fully agree.

But the similarities and differences between the systems of Keynes-Shackle and RCO are too manifold and complex to analyze in this already overloaded article; that will be done separately. Meanwhile, I still wish to alert the readers to this extremely important possibility. Therefore I mention it in the introductory summary of the potential benefits of RCO – in a qualifying context until proven true.

Of course, this does not mean at all that RCO is in the same category as the ideas and systems of these two giants. RCO is just a tool. But in the field of decision-making, no system is complete without realistic and effective means of its implementation.

The main inspiration for RCO was the seminal book of Luce and Raiffa (Luce and Raiffa, 1957). As mentioned in Section 3, the book showed that none of the strategy comparison criteria was “the best,” and none could be used as a solitary guide. Obviously, that implicitly indicated the need for developing multiple strategies and multiple scenarios, for screening strategies by looking at their payoffs under several criteria over the whole range of scenarios, and then choosing a strategy subjectively from the remaining subset of strategies.

This approach, however, could be implemented in two ways. The first, “immediate gratification” approach, is simply to follow this advice to the letter:

- Construct a number of single-scenario “What if” linear programming models and solve these models;
- Cluster the groups of similar model solutions into a number of strategies;
- Construct single-scenario linear programming models for all “strategy vs. scenario” combinations and solve these models;
- Use the obtained values of the objective function as the needed payoffs;
- “Weakly” screen the strategies by applying to payoffs several criteria.

In the 1960s, a group of scientists from the USSR energy industry, also under the influence of Luce and Raiffa, used this approach to develop a scenario planning system called “the Zone of Uncertainty,” where the “zone” was the screened subset of strategies. The system was described by Makarov and Melentyev (Makarov and Melentyev, 1973, pp. 115–257).

For its time, this approach proved enormously progressive. The methodology included an excellent mathematical description of the behavior of dynamic energy systems, as well as other novelties. Neither of the later RCO inventions (namely, the multiscenario models, the outcome matrix, the risk-limiting constraints, “strong” screening, the utility conversion methods, the new “synthetic criteria,” and the “strategic frontiers”) was yet used. Nevertheless, strategies were developed, checked against up to 1500 scenarios, and “weakly” screened by several then existing criteria. The results were good, robust compromises between scenario extremes in using different fuels or regions. In 1970, many of the derived results were used in the plans for the development of the USSR energy industry (Makarov

and Melentyev, 1973, pp. 115–257). As far as I know, the methodology is not in use now.

RCO took from that approach two important ideas: first, splitting the variables of mathematical programming problems into “strategic” and “operational” groups, as described in Section 5(iii), although it was used there in single-scenario, rather than in multisenario, models; second, using the scenario outcomes not of strategies proper, but of their contingency plans.

The less obvious second way requires not one, as in “the Zone of Uncertainty,” but several layers of protective filters – including the “strong” filters – and selecting the strategy only after applying all of them. Protection from risk is performed here at several consecutive stages of the system and becomes dominant and comprehensive. In other words, this approach requires the full range of novel techniques later introduced in RCO (see above). Each of these mutually reinforcing techniques substantially improves the overall reliability of the system. Naturally, this more rewarding way, chosen for RCO, is also more difficult and requires much longer time to develop.

9 An example

It is evident that RCO is most useful in situations of great uncertainty, especially in long-range planning. *Moreover, its flexibility allows applying RCO in quite realistic and therefore very complex problems.* We will, however, demonstrate an application of some techniques of RCO on a very simple short-term problem.

A furniture company manufactures and sells chairs, tables, and beds. Net profit from the sale of one chair is \$3, one table \$5, and one bed \$3.50. Manufacturing a chair requires 0.15, a table 0.45, and a bed 0.25 units of plywood. The company needs a special type of plywood; 45 units of it are available. The company expects the market to absorb up to 160 chairs, 80 tables, and 120 beds. The company wants to find out how many pieces of each type of furniture should be manufactured in the planning period to maximize its profit, while keeping its share in the market of chairs at no less than 45 percent. Production “for the shelf” is forbidden.

The problem can be described by a linear programming model that has an objective function (to maximize profit), seven explicit constraints (in addition to implicit constraints of non-negativity for all variables), five variables, and 11 input parameters. As mentioned in Section 5, RCO uses “bookkeeping” variables. In this model, two “bookkeeping” variables are the total profit and the market share in the chair market, while three other variables are production of chairs, tables, and beds. The 11 input parameters are: profit per unit of furniture, plywood requirement per unit of furniture, market demand for furniture (each of these three parameters for three types of furniture), the availability of plywood, and the minimally acceptable market share in chairs. The seven constraints are: definition of the total profit through the production volumes of the three types of furniture; total availability of plywood (defined through the same production volumes); upper bounds on these three production volumes, defined by the capacity of each market; definition of the market share in chairs through the chair production and chair market capacity; and the lower bound on that market share.

Table 1. Scenarios

Scenario	Weight	Demand for chairs	Net profit for a bed (\$)
1	0.09	100	2.50
2	0.06	100	5.00
3	0.36	150	2.50
4	0.24	150	5.00
5	0.15	220	2.50
6	0.10	220	5.00

The optimal solution of this single-scenario model corresponds to production of 160 chairs, no tables, and 84 beds. The total profit equals \$774, and the market share in chairs is 100 percent.

However, the input values have been obtained from forecasts and “guesstimates.” The sources of information are not sufficiently reliable, so that the real future values of some input parameters may be different. That puts the company at risk of not meeting its financial and market share expectations. To protect itself against such risks, the company applies RCO. The goal is to develop a limited-risk strategy that is as good as possible.

Let us suppose that all but two of the input values are considered sufficiently reliable. The two “uncertain” parameters are: potential demand for chairs, which may vary between 100 and 220 pieces, and net profit for a bed, to vary between \$2 and \$5. The company decides to use three alternative values of demand for chairs (100, 150, and 220 pieces) and two potential values of the net profit for a bed (\$2.50 and \$5).

The company has some data about the probabilities of the possible alternatives. They are unreliable, but still can be used. They are: for chairs 0.15, 0.60, and 0.25, and for beds 0.60 and 0.40. Under these probabilities, the expected demand for chairs and the expected profit for beds indeed are 160 and \$3.50, as assumed in the deterministic model.

These alternatives generate $3 \times 2 = 6$ possible combinations, or scenarios, shown in Table 1. A scenario is a combination of values for all 11 parameters, with identical values under all scenarios only for 9 “certain” parameters. The weight (probability) of each scenario equals the product of probabilities of the parameter value alternatives that generate the scenario. For instance, the weight of Scenario 1 equals $0.15 \times 0.60 = 0.09$. For simplicity, we do not correlate the net profit and demand alternatives.

As noted in Section 5, we start with formulating and solving six mutually independent “what if” linear programming models for six scenarios. (For instance, in the model for Scenario 6 the demand for chairs is 220 and the net profit for a bed is \$5.) The derived strategies are dissimilar, so we proceed to the extended RCO.

Now we formulate a six-scenario model that provides a solution optimal over all scenarios combined. In essence, the model builds six single-scenario submodels, similar to six “what if” models, and then holds them together – first, by the overall objective function, and, second, by assuring that the value of each “strategic variable” does not change (or changes within specified limits) from one scenario

Table 2. Strategy A

Scenario	Weight	Profit	Chairs	Market share	Tables	Beds
1	0.0900	633	100	1.00	67	0
2	0.0600	900	100	1.00	0	120
3	0.3600	633	100	0.67	67	0
4	0.2400	900	100	0.67	0	120
5	0.1500	633	100	0.45	67	0
6	0.1000	900	100	0.45	0	120
Total/avg.	1.0000	740	100	0.66	40	48

Table 3. Strategy B

Scenario	Weight	Profit	Chairs	Market share	Tables	Beds
1	0.0900	620	100	1.00	40	48
2	0.0600	740	100	1.00	40	48
3	0.3600	620	100	0.67	40	48
4	0.2400	740	100	0.67	40	48
5	0.1500	620	100	0.45	40	48
6	0.1000	740	100	0.45	40	48
Total/avg.	1.0000	668	100	0.66	40	48

to another. Let us consider the target production of chairs a strategic variable; then this target will be met under every scenario, whatever that requires and whatever may be the market conditions.

Since, under Scenarios 5 and 6, the production of chairs should be no less than $220 \cdot 0.4 = 99$, let us fix the value of that strategic variable at 100 pieces. The optimal solution for this model is shown in Table 2. The resulting initial strategy will be called Strategy A. This optimal solution clearly demonstrates the risks involved in a very simple and realistic application of optimization models, even under the short-term planning. The assumed range of price and demand alternatives leads to wild swings in production of tables and beds: whenever the profit for beds is high, all disposable resources of plywood are used to manufacture beds; otherwise, tables. Because of the instability of its “all-or-nothing” solutions, the optimization model introduces here its own risk component. (In the last row, “Total/avg.” means “total” for the “Weight” column and “average” for other columns.)

We find out now that our model is incomplete; it does not take into account the need for stability of production and employment. Can the company fix the production targets for tables and beds – say, at the average levels of Strategy A – to smooth its activities over different scenarios of the future? When these targets are fixed, the company’s profit is a respectable \$668. This certainly is a possible course of action; it will be called Strategy B and is shown in Table 3.

The problem has five interconnected activities, specifically, production of three types of furniture, “getting profit,” and “getting market share.” The last two can be considered as the risk types, since the management is concerned with the level of these activities. The company would like to watch over and, as far as possible, to

Table 4. Profit

Scenario	Weight	Strategy A	Strategy B
1	0.0900	633	620
2	0.0600	900	740
3	0.3600	633	620
4	0.2400	900	740
5	0.1500	633	620
6	0.1000	900	740
Total/avg.	1.0000	740	668

Table 5. Market share

Scenarios	Weight	Strategy A	Strategy B
1 & 2	0.1500	1.00	1.00
3 & 4	0.6000	1.00	0.67
5 & 6	0.2500	0.68	0.45
Total/avg.	1.0000	0.92	0.66

control the outcomes in these two activities. For that purpose, RCO constructs the outcome matrix. This three-dimensional matrix consists of two two-dimensional tables “scenario vs. strategy,” each for one type of activity, Table 4 for profit and Table 5 for the market share. In Table 5, the probability of a lower bed profit does not matter; therefore, weights for Scenarios 1 and 2 can be considered jointly, as well as 3 and 4 or 5 and 6. Strategy A, which is unacceptable because of the lack of stability, is included in these tables just for the sake of comparison.

Suppose now we want to discard the condition that, under all scenarios, the production of tables and beds equals the averages, that is, 40 and 48, respectively. Instead we impose a weaker, risk-limiting “80 percent stability” constraint: production of tables under Scenarios 2, 4, and 6 should be no less than 0.8 of table production under Scenarios 1, 3 and 5, respectively.

It turns out that, depending on the values of bed profit probabilities, this model has two optimal solutions. Suppose that we are relatively optimistic and believe that the probability of the low (\$2.50) net profit for a bed is indeed 0.60 or somewhat higher, say, in a “probability range 1” from 0.6000 to 0.8648 (the last number is explained below). The mid-range probability for range 1 equals 0.7324. Let us find the optimal strategy for that value of the bed profit probability, changing the scenario weights correspondingly. This optimal solution (to be called Strategy C) is shown in Table 6.

This solution suggests that, to maintain “80 percent stability,” one line of furniture should be abandoned and manufacturing should be concentrated on two types, chairs and beds. The solution has certain advantages in comparison with Strategy B, including higher profits under three “fat” scenarios (\$900 instead of \$740), but it also has lower profits under “lean” scenarios: \$600 instead of \$620. The average profit is now \$680. (Under the new scenario weights, the average profit of Strategy B is no longer \$668; it decreases to \$652.)

Table 6. Strategy C. In “Probability Range 1”

Scenario	Weight	Profit	Chairs	Market share	Tables	Beds
1	0.1099	600	100	1.00	0	120
2	0.0401	900	100	1.00	0	120
3	0.4395	600	100	0.67	0	120
4	0.1605	900	100	0.67	0	120
5	0.1831	600	100	0.45	0	120
6	0.0669	900	100	0.45	0	120
Total/avg.	1.0000	680	100	0.66	0	120

Table 7. Strategy D. In “Probability Range 2”

Scenario	Weight	Profit	Chairs	Market share	Tables	Beds
1	0.1399	633	100	1.00	66	1
2	0.0101	685	100	1.00	53	24
3	0.5595	633	100	0.67	66	1
4	0.0405	685	100	0.67	53	24
5	0.2331	633	100	0.45	66	1
6	0.0169	685	100	0.45	53	24
Total/avg.	1.0000	680	100	0.66	65	3

However, if we are more pessimistic and consider that the probability of the low bed profit is 0.8649 to 1.0, or in a “probability range 2,” the model provides another solution (Strategy D), in which both tables and beds are manufactured. Strategy D meets the “80 percent stability” condition in a different way, as shown in Table 7. (We required stability only for tables, not for beds.) This solution is found for the “mid-range 2” probability of 0.9324, with the corresponding scenario weights.

In Strategy B, we produce 100 chairs, 48 beds, and 40 tables under every scenario; in Strategy C, the corresponding levels are 100, 120, and 0; in Strategy D, the production levels vary within the specified limits. Tables 4 and 5 should now be expanded into Tables 8 and 9. Instead, Table 8 has to be split into Tables 8a and 8b, because, due to the changing scenario weights, the average profits for Strategy D differ for two ranges of bed profit probabilities. We exclude Strategy A from all tables.

In general, Table 9 should also be split among the different probability ranges, but, as indicated earlier, in this example the market share of chairs does not depend upon the probability values for beds.

To demonstrate the validity of statements made in Section 6, let us carry this analysis a little further and, once again, compare the strategy profits for the mid-range values of “range 1” and “range 2.” As shown in Tables 8a and 8b, the optimization (linear programming) models indeed provide here “all-or-nothing,” extremal solutions, easily switching from one extreme to another. In “range 1,” Strategy C, which completely excludes the production of tables, seems to be better than the “compromise” Strategy B. In “range 2,” Strategy D is the best; it almost completely

Table 8a. Profit. In Probability Range 1

Scenario	Weight	Strategy B	Strategy C	Strategy D
1	0.1099	620	600	633
2	0.0401	740	900	685
3	0.4395	620	600	633
4	0.1605	740	900	685
5	0.1831	620	600	633
6	0.0669	740	900	685
Total/avg.	1.0000	652	680	647

Table 8b. Profit. In Probability Range 2

Scenario	Weight	Strategy B	Strategy C	Strategy D
1	0.1399	620	600	633
2	0.0101	740	900	685
3	0.5595	620	600	633
4	0.0405	740	900	685
5	0.2331	620	600	633
6	0.0169	740	900	685
Total/avg.	1.0000	628	620	637

Table 9. Market share

Scenarios	Weight	Strategy B	Strategy C	Strategy D
1 & 2	0.1500	1.00	1.00	1.00
3 & 4	0.6000	0.67	0.67	0.67
5 & 6	0.2500	0.45	0.45	0.45
Total/avg.	1.0000	0.66	0.66	0.66

excludes beds under the depressed bed profits, and, under the higher bed profits, produces tables only as far as it is required by the “stability constraint.”

Note, however, that this analysis involves only the issue of which candidate strategies are considered “the best” by the optimization models. It says nothing about which candidate should be selected. As will be seen below, “synthetic criteria” use different argumentation.

Now, matrices 8a and 9 make up the three-dimensional “outcome matrix” for “range 1,” while 8b and 9 comprise the similar matrix for “range 2.” Taken together, the two three-dimensional matrices form a four-dimensional matrix “scenarios vs. strategies vs. risk types vs. probability ranges” for this example.

Imposing other risk-limiting constraints would lead to other candidate strategies with different trade-offs and risk/return characteristics. For instance, we could request “100 percent stability.” In that case, there again exist two optimal solutions for different probabilities, but “range 1” would be a little wider – from 0.6000 to

Table 10. Generalized profit

Scenario	Weight	Strategy B	Strategy C	Strategy D
1	0.1099	630	610	643
2	0.0401	750	910	695
3	0.4395	627	607	640
4	0.1605	747	907	692
5	0.1831	624	604	637
6	0.0669	744	904	689
Total/avg.	1.0000	659	687	654
Best	–	750	910	695
Worst	–	624	604	637

0.8888, instead of 0.8648. Also, only “all-or-nothing” are optimal: either tables or beds are completely excluded from the optimal solution.

But we will not further expand this example and will, instead, emphasize the following important feature of the “analysis of shifting probabilities” demonstrated above: it helps clarify our perception of what is “optimism” and “pessimism” in this problem. That, in turn, helps at stage (ix) of RCO, when we compare strategies by means of applying the “synthetic” criteria in the framework of the strategic frontier.

If we consider the lower “range 1” of depressed bed profits, then we are more optimistic, and vice versa. Strategy D seems to be much worse than Strategy C: in the “fat” scenarios, the D profit is \$685 instead of \$900, which is not compensated by a small increase of \$33 (\$633 instead of \$600) in the “lean” scenarios, in spite of their higher weights. To switch into “range 2” which gives preference to Strategy D, we must have solid grounds to be very pessimistic. Let us, therefore, be cautiously optimistic, so that we can justifiably proceed with only one three-dimensional matrix that consists of Tables 8a and 9.

As mentioned earlier, for comparing strategies, we need to convert a now three-dimensional “outcome matrix” into a two-dimensional payoff matrix, that is, to convert the market share to profit. For the sake of simplicity, we do not use the technique of multiple conversion coefficients stated in Section 5(vii) and assume, instead, that a single value of the market share conversion coefficient is sufficient. Suppose that 0.01 of the market share for chairs corresponds to \$0.10 of profit. The resulting payoff matrix for the “generalized profit” is shown in Table 10. For comparing strategies, we will need to know the profits for the “best case” and the “worst case,” as well as the average profits; we therefore present all these numbers.

Note that if we had two Tables 10 for “range 1” and “range 2,” they would differ only in the “Weight” column and the “Average” row. The “Best” and “Worst” rows would be identical.

We will compare the strategies by two “synthetic” criteria, the “PO index” and the “PI index” mentioned in Section 5(viii).

For the “PO index” criterion shown in Figure 2, we need two numbers for each strategy: the “best case” and the “worst case” payoffs. As follows from Table 10, for Strategy B these numbers are \$750 and \$624, for Strategy C – \$910 and \$604,

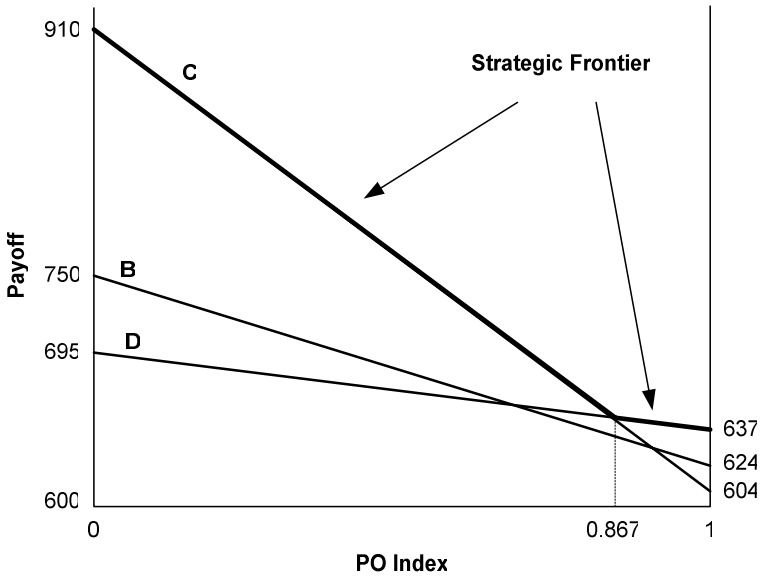


Fig. 2. “Pessimism – optimism index” criterion

and for Strategy D – \$695 and \$637. The zero value of the PO index corresponds to complete optimism; the 1.0 value, to complete pessimism. For each strategy, the “best case” value is plotted on the left vertical axis, the “worst case” value, on the right vertical axis. These two points are then connected by a straight line. The upper segments of these straight lines form the “strategic frontier.”

According to Figure 2, this frontier lies with Strategy C at the index values from 0 to 0.867, and with Strategy D, from 0.867 to 1. This means that Strategy D is preferred only in a small, most pessimistic portion of the possible index values. Since we are relatively optimistic, Strategy C is preferable. As may be expected, the “compromise” Strategy B is worse at the whole [0, 1] segment of the PO index values.

The strategic frontier of Figure 2 has an important distinction from the “efficient frontier” of portfolio theory. The strategies that form the strategic frontier move from left to right, from more optimistic to more pessimistic. Since the more optimistic strategies start higher and decline more steeply (otherwise they would not intersect with the next strategy), the strategic frontier – in contrast to the “efficient frontier” – is always concave upward.

The PO index has many advantages, the main one being complete independence from the scenario probabilities; instead of “probability,” it explores “possibility.” (In this example, the picture in Fig. 2 does not change from one probability range to another.)

However, this criterion might be too optimistic, especially if, as in this example, we did not take correlation into account: the very best scenario may be difficult to achieve. The second criterion, the “PI index,” assumes, instead, that the best that we can expect is not the “best case” payoff but, rather, the weighted average payoff.

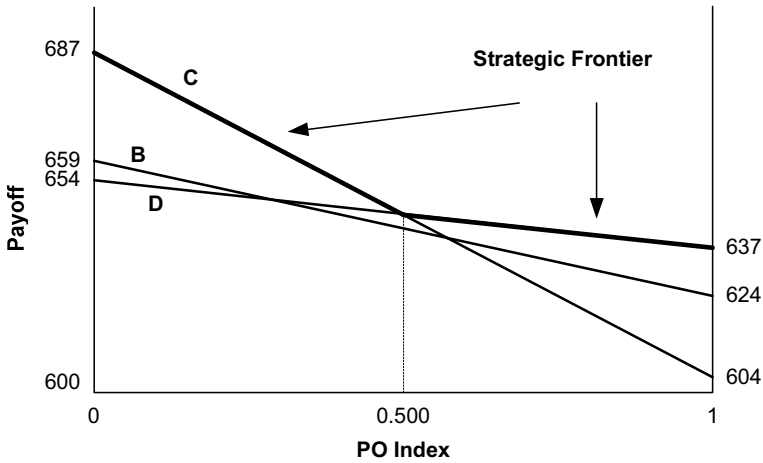


Fig. 3. "Partial ignorance index" criterion

This value (see Fig. 3) is plotted on the left vertical axis, while the "worst case" value is plotted on the right axis. Otherwise, these two criteria are similar. For each strategy, we again need two numbers: the weighted average and the "worst case" payoff.

For Strategy B, they are \$659 and \$624, for Strategy C – \$687 and \$604, and for Strategy D – \$654 and \$637. According to Figure 3, Strategy D is preferable now on the broader interval of the PI index values – from 0.5 to 1.

But this is a maximization point of view, which does not necessarily prevail in Decision Science, where the Tortoise is often preferred to the Hare. Indeed, if we are either cautious optimists or hopeful pessimists, our value of the PO or PI index would perhaps be somewhere in the neighborhood of the [0.5, 0.867] interval. In Figure 2, at the 0.867 value of the PO index, the weighted profit of Strategy B is \$641, or only \$4 less than that of both Strategy C and Strategy D. In the more cautious and probably more realistic "PI index" graph of Figure 3, the expected profit of Strategy B is contained within being worse than Strategy C by \$28 (at the index value of 0.0) or worse than Strategy D by \$13 (at the index value of 1.0). At the inflection point of the "strategic frontier," that is, at the 0.5 value of the index, the expected profit of B is \$642, or only \$4 less than that of the two other strategies.

As shown above, strategy D is much worse than C. In production stability, D is worse than both C and B. So we decide to discard D and consider the two strategies, B and C. Each has 100 percent stable production levels for all three types of furniture; B allows the company to stay in both table and bed markets, while C abandons the tables. The decision-maker would judge what is better; additional considerations may be needed.

These results are, to a certain degree, predictable, but, again, our goal was to demonstrate the techniques effectively, especially the possibility of narrowing our implied range of the optimism index.

The decision-maker will choose the strategy subjectively, but, we hope, by taking into account all these analyses and conclusions. In the small-scale and simple

problem of this example, such detailed analysis might seem excessive, but in a large-scale strategic management problem it might save us from serious mistakes. It is better to be safe than sorry.

10 Conclusions

1. In the current fragile geopolitical and economic environment, when the probable course of future events is highly uncertain, current methods and models of business decision-making in strategic management may lead to serious mistakes. To upgrade the quality of the process, a new prescriptive approach is proposed. It uses an ensemble of special models, procedures, and algorithms to generate, evaluate, and help in the execution of good, robust, and limited-risk strategies. The ensemble still uses maximization methods, but only in combination with several protective filters that screen, modify, and scale back the strategies, as necessary. It has to meet a number of rigorous requirements. The role of this ensemble is critical: as with every type of protective equipment, it replaces, to some extent, the needed knowledge about the future.

2. Risk-Constrained Optimization[®] (RCO), a powerful system that has been granted a US patent, contains this ensemble and fully meets demanding requirements. RCO is a system of planning under uncertainty that searches for the most acceptable compromise between improving the results and reducing the risk in our decisions.

3. RCO is comprised of three major parts: (a) development of a large number of scenarios; (b) generation of a number of candidate strategies that are good, flexible, robust, and risk-limited; and (c) identification of a few of the most suitable candidates and final subjective selection of the strategy. Its two main ideas are: first, that the formal methods not reliable enough to find the best alternative are still good enough for an easier task – to screen out the worst; second, that the decision-maker has to be actively involved in the whole process of generation, evaluation, and screening of alternative strategies.

4. RCO replicates the “natural” process of decision-making (“multiple everything” – scenarios, strategies, goals and risk types, with simplified selection of the preferable strategy), not deviating from it by a single unwarranted assumption. More sophisticated utility-maximization decision methods are considered just a special case of the “natural” process. Accordingly, these methods are relegated from the general, top-level conceptual paradigm of the final strategy selection to auxiliary operations, which may, or may not, be used at the earlier stages of constructing and screening strategies.

5. RCO is an honest and hard-to-manipulate system. It also offers several important advantages. The RCO system:

- For the first time in more than 50 years, legitimates the high-level use of both computers and optimization models in strategic management.
- Creates a protective device that is most reliable, since it screens the strategies by tools unsurpassed in their meticulousness.

- Allows transforming standard tools into realistic customized models, even by executives who have no modeling skills.
- Might provide the tools for implementing the uncertainty-centered decision systems of J. M. Keynes and G. L. S. Shackle.

Why should so many diverse breakthroughs come from one system? It is, I think, because RCO is “a natural,” further strengthened through the use of computers and advanced computational methods. Results come easy to “a natural.”

6. RCO might become the “disruptive technology” so sorely needed to replace the current, greatly deficient approaches to strategic decision-making.

References

- Bossidy L, Charan R (2002) *Execution: the discipline of getting things done*. Crown Business: New York
- DeLoach JW (2000) *Enterprise-wide risk management: strategies for linking risk and opportunity*. Prentice Hall: London New York
- Donaldson G, Lorsch JW (1983) *Decision making at the top. The shaping of strategic direction*. Basic Books: New York
- Fleming I. (1959) *Goldfinger*. Macmillan: New York
- Hoch SJ et al. (2001) *Wharton on making decisions*. Wiley: New York
- Hurwicz L (1951) *Optimality criteria for decision making under ignorance*. Cowles Commission Discussion Papers, Statistics, No 370
- Jehn KA, Weigelt K (2001) *Reflective versus expedient decision making: views from East and West*. In: Hoch SJ et al. (2001), pp 103–114
- Kleindorfer PR et al. (1993) *Decision sciences*. Wiley: New York
- Knight FH (1921) *Risk, uncertainty and profit*. University of Chicago, Chicago
- London S (2003) *Thinking ahead is just the easy part*. Financial Times, July 14
- Lovallo D, Kahneman D (2003) *Delusions of success: how optimism undermines executives' decisions*. Harvard Business Review 81(7): 56–63
- Luce RD, Raiffa H (1957) *Games and decisions*. Wiley: New York
- Makarov AA, Melentyev LA (1973) *Methods of studying and optimizing the energy sector*. Nauka: Novosibirsk, USSR
- Miccolis JA et al. (2001) *Enterprise risk management: trends and emerging practices*. The Institute Of Internal Auditors Research Foundation, Altamonte Springs, FL
- Radner R, Shepp LA (1996) *Risk vs. profit-potential: a model for corporate strategy*. Journal of Economic Dynamics and Control 20: 1373–1393
- RiskMetrics™ – Technical Document (1996) Morgan – Reuters: New York
- Rubin R, Weisberg J (2003) *In an uncertain world. Tough choices from Wall Street to Washington*. Random House: New York
- Senge PM (1990) *The fifth discipline*. Doubleday Currency: New York
- Simon HA (1982) *Models of bounded rationality*. MIT Press: Cambridge, MA
- Singhal V, Hendricks K (2002) *How supply chain glitches torpedo shareholder value*. Supply Chain Management Review 6: 18–24
- US Patent (1999) # 5,930,762 of July 17
- Voyle S (2002) *Lessons from the man in the street*. Financial Times, December 16