

System Dynamics and Decisions Under Uncertainty

Erling Moxnes¹
Chr. Michelsen Institute
Fantoftvn.38, N-5036 Fantoft, Norway

Abstract

The theory of decisions under uncertainty share basic assumptions with system dynamics. Both methods require that decisions are based on only available information, and both methods focus on the development of policy rules that improve system performance. Both methods have other implications for parameter estimation than conventional deterministic analysis. Fluctuations are frequently studied in system dynamics, and fluctuations and randomness are of great importance for decisions under uncertainty. Decisions under uncertainty can be studied by analytical methods, dynamic programming and Monte Carlo simulations. The latter method is quite easily applied to system dynamics models. Using Monte Carlo simulations we show that uncertainty has important implications for decisions influencing the "greenhouse" effect. Note that risk aversion is not an issue in this example. The theory of decisions under uncertainty brings new qualitative insights to system dynamics, and facilitates quantitative improvements of policy rules. Referring to or applying the theory of decisions under uncertainty might help to get a wider academic acceptance of system dynamics models, which are often thought of as being realistic but quite uncertain. The principles of system dynamics might bring the field of decisions under uncertainty in the direction of greater realism. The focus on real life interpretation of system dynamics models is most useful for the application of apriori information. Apriori information is needed to establish important autocorrelation in cases where short time-series do not contain sufficient information.

System Dynamics and Decisions Under Uncertainty

Optimization has not got a strong foothold within the field of system dynamics, particularly when dealing with uncertain social systems. According to R.F.Naill: "...uncertainties and imprecisions inherent in social systems modeling make any claim of an 'optimal' [deterministic] policy design unjustified", (Naill 1974, 16). This is also what scientists within the field of stochastic optimization or decisions under uncertainty argue about deterministic optimization. According to C.Henry: "...in general, replacing random variables by their expectations will not lead to the appropriate decisions" (Henry 1974, 1007).

To my knowledge the technique and philosophy of stochastic optimization or decisions under uncertainty have not been explicitly considered in the system dynamics literature. The purpose of this paper is to demonstrate that these two fields share underlying assumptions. Both methods require that decisions are based on only available information, and both methods focus on the development

of policy rules to improve system performance. The purpose is also to show that the methods might empower each other. For system dynamics, stochastic optimization might provide both qualitative insight as well as better quantitative policy suggestions. For decisions under uncertainty, the principles of system dynamics and its focus on real life interpretation and a priori data, might provide better models and understanding of random processes. The paper presents a Monte Carlo method for making decisions under uncertainty that is easily applicable for system dynamics models.

The paper is organized as follows. First, an example is used to introduce the problems of decisions under uncertainty. Secondly, the Monte Carlo method is presented and used to find an optimal strategy in case of an irreversible decision. Thirdly, various stochastic optimization methods and modeling techniques are discussed.

Examples of decisions under uncertainty

The general problem of decisions under uncertainty can be stated as follows:

$$\max_x E \{ U(x, \varepsilon) \} \quad (1)$$

In addition to the ordinary problem of formulating and maximizing a criterion $U(x)$, uncertainty introduces three new aspects:

- Uncertainty represented by the random variable ε must be formulated and estimated.
- The expected value of the criterion $E \{ U \}$ must be found.
- The optimal solution must not assume the use of information that is not yet available ("You know that you'll know, but you don't know yet").

Risk aversion can, but does not have to be part of the problem. (It can be built into the criterion U).




Figure 1: Examples of random and deterministic fluctuations in oil and gas prices.

An example of a decision under uncertainty is the choice between investing in a power station that uses oil or a multi-fired one that can switch between using oil and gas. If future oil and gas prices are expected to be exactly equal, profits are maximized by choosing the oil-fired station which is cheaper than the multi-fired station. If there are independent random fluctuations in the two prices, while the average or expected values are still the same, the multi-fired station becomes more valuable because it benefits from utilizing the cheaper fuel all the time, see figure 1. Whether the multi-

fired station should be chosen depends on the amount of randomness. Thus, an estimate of the randomness is needed to make the correct decision.

If there is a certain adjustment time in going from one fuel to another, for example due to delivery delays of the fuels, it will not be possible to use the cheapest fuel all the time. The multi-fired option will only pay off each time the gas price is lower than the oil price for a longer period than the adjustment time. Thus, adjustment delays reduces the value of the multi-fired option. On the other hand, if price fluctuations are deterministic, like in the second exhibition of figure 1, it would be possible to plan ahead and to get the full benefit of multi-firing in spite of adjustment delays. In between these two extreme cases is the more realistic case with autocorrelated randomness.

To conclude, randomness and autocorrelation are important for the correct choice between an oil-fired and a multi-fired power station. This conclusion is independent of whether the criterion is measured in money or in units of utility. The purpose of the decision under uncertainty is to maximize the expected outcome.

If the criterion is measured in units of money, we see that stochastic optimization increases the expected value, while an insurance policy would reduce the expected value.

Note that decisions under uncertainty bring a new meaning to forecasting. Decisions under uncertainty acknowledge the scepticism of (Forrester 1961, 431): "If the presence of noise is admitted, we must necessarily come to the conclusion that even the perfect model may not be a useful predictor of the specific future state of the system it represents." However, decisions under uncertainty go a step further and take advantage of the fact that scenarios of the future might belong to the same "family" as Forrester terms similarities with respect to amplitudes of excursion and abruptness of change. In Henrik Ibsen's terminology in his play Hedda Gabler: "Jørgen Tessman: About the future! My God, we know nothing about it. Eilert Løvborg: No, but there are still things to be said about it."

Stochastic methods have been applied in studies of optimal stopping (e.g. job search), futures pricing, option pricing, portfolio selection, asset pricing, project evaluation etc. (Malliaris and Brock 1987) refer numerous examples.

Monte Carlo simulations to make decisions under uncertainty

In this section an example demonstrates the Monte Carlo² method used to make decisions under uncertainty. The next sections will explore details of the methodology and compare it to other techniques.

In short, the method goes like this. Propose a rule-of-thumb (heuristics) for decisions based on available information as this information is revealed over time. Start simulating. Draw random numbers from specified distributions to mimic the revelation of information. Make decisions based on this information, and evaluate the criterion as the results appear. Repeat the simulations a large number of times, and calculate the expected (average) value of the criterion. Repeat this whole process with various rules-of-thumb, and choose the rule that maximizes the expected criterion.

The following practical example is inspired by (Henry 1974), who has studied the effect of irreversibility on an investment decision under uncertainty. Think of our irreversible decision as being releases of CO₂ into the atmosphere³. Think of the investment decision as being a power station burning coal and releasing CO₂. The benefits of power generation are assumed to be known precisely, while the costs of releasing CO₂ are uncertain. However, measurements of the effects of historical and on-going releases of CO₂ as well as basic research, will provide new information about the costs over time.

The following formulation is used to represent the stochastic process of revealed expected costs $\varepsilon(t)$:

$$\varepsilon(t) = \int_{\tau=0}^t v(\tau) d\tau + \varepsilon_0 \quad (2)$$

The expected costs of having the extra units of accumulated CO₂ from the power plant equals ε_0 initially. This value must be exclusively based on apriori knowledge about the problem, since significant changes in climate and costs have not been measured yet. At each point in time new information is revealed through the random variable $v(t)$, which is a uniformly distributed variable. When simulating, we assume that $v(t)$ is updated every fifth year (simple Euler integration with a five year step). Initially it ranges from -0.5 to +0.5 per five years. Over time this range decreases linearly towards no variation in year 150. Thus, we assume that after 150 years there is absolute certainty about the costs of CO₂ releases⁴.

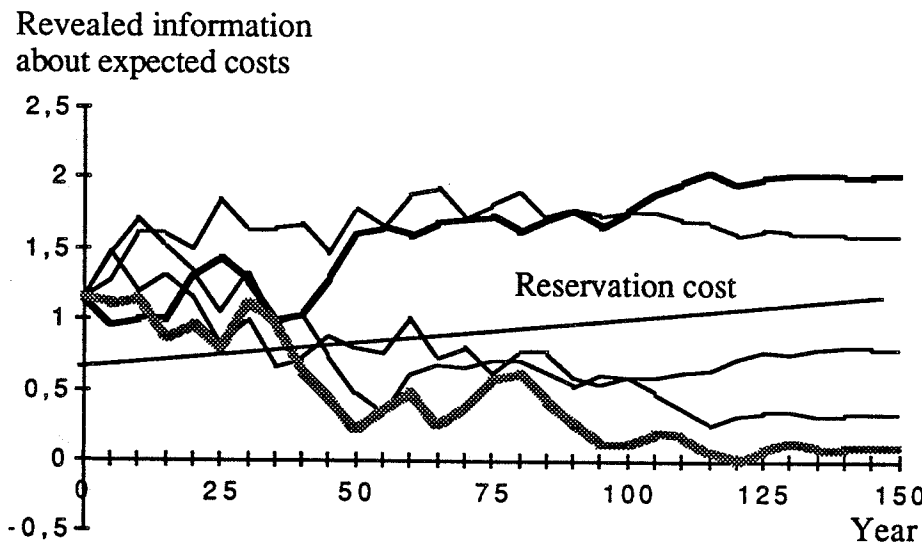


Figure 2: 5 runs of a Monte Carlo simulation of revealed expected costs of CO₂ releases.

Figure 2 shows five runs of a Monte Carlo simulation of how the revealed expected costs develop. Notice how variation from year to year declines over time until full certainty is achieved in the end of each scenario. The parameters of the random process are set such that the distribution of the costs in the final year correspond to the distribution of today's apriori cost estimates. We notice that the expected value in year 150 is the initial expectation, and the standard deviation is 0.82⁵. The stochastic process is similar to a random walk with the exception that $v(t)$ is uniformly instead of normally distributed and the variance declines over time.

The rule-of-thumb for the investment decision is such that the investment will not take place unless revealed expected costs get below a reservation cost ε_r . The first time $\varepsilon(t)$ gets below the reserva-

tion cost, the "decision time" t_d is set equal to t . The variable $d(t)$ denotes that an irreversible decision has been made by switching from 0 to 1. The formal rule is

$$t_d = \begin{cases} 150 & \text{initially} \\ t & \text{at the time } \varepsilon(t) \leq \varepsilon_r \text{ for the first time} \end{cases}$$

$$d(t) = \begin{cases} 0 & \text{if } t < t_d \\ 1 & \text{if } t \geq t_d \end{cases} \quad (3)$$

The benefits from the power station $b(t)$ are set equal to 1 by definition, during the 30 year period the station operates. Otherwise there are no benefits.

$$b(t) = \begin{cases} 1 & \text{if } t_d \leq t < t_d + 30 \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

We assume that there is a 50 years delay T from the power plant starts operating ($d(t)=1$) to the full effect of the costs are felt in terms of climate effects⁶. The distribution of the costs over time $c(t)$ is given by a first order delay:

$$c(t) = \int_{\tau=0}^t (d(\tau) - c(\tau)) / T \, dt + 0 \quad (5)$$

We choose the net present value (*NPV*) of benefits $b(t)$ minus costs $c(t) \cdot \varepsilon_{150}$ to be the criterion. Notice that ε_{150} represents the true, revealed costs after 150 years. The exponent takes care of the discounting with a discount rate $r=0.02$ per year⁷. The time horizon is 150 years, which should give a sufficient accuracy with the chosen discount rate⁸.

$$NPV = \int_{t=0}^{150} e^{-rt} (b(t) - c(t) \cdot \varepsilon_{150}) \, dt \quad (6)$$

We now set ε_0 and consequently the expected value of ε_{150} equal to 1.148, such that the expected $NPV=0$ when the power station is started in the initial year. This enables us to focus on the net effect on the NPV of using a better decision rule than the immediate production start. In other words, ε_0 is set such that the traditional deterministic analysis based on expected costs, yields a net present value of 0.

The reservation cost ε_r in the rule-of-thumb is assumed to increase over time, to take account of reduced uncertainty:

$$\varepsilon_r = \varepsilon_{r0} \left(1 - \frac{t}{150}\right) + \varepsilon_0 \cdot \frac{t}{150} \quad (7)$$

When $t=150$, there is no uncertainty, and consequently the power station should be built if the NPV is positive. Thus, at $t=150$ the reservation cost $\varepsilon_r = \varepsilon_0$. Initially, uncertainty is large, and by using a reservation cost ε_{r0} lower than ε_0 , the project start is delayed to wait for more information. The reservation cost increases linearly between start and end. The purpose of the Monte Carlo simulation is to establish the value of ε_{r0} that gives the maximum expected NPV .

The expected NPV is calculated as the average of the individual net present values NPV_i of n runs of a Monte Carlo simulation:

$$E(NPV) = \frac{1}{n} \sum_{i=1}^n NPV_i \quad (8)$$

This concludes the presentation of the model.

Then we are ready to search for the reservation cost ε_{r0} that maximizes the expected NPV . We enumerate with ε_{r0} ranging from 0.2 to 0.9999 relative to ε_0 . Table 1 shows the resulting expected NPV 's with standard deviations, as well as average investment years and number of cases when power plants are built. Remembering that the expected NPV in the case with immediate production start ($\varepsilon_{r0}/\varepsilon_0 \geq 1$) is zero, we see that all the suggested values of ε_{r0} improves the expected net present value. Clearly, using the procedures for decisions under uncertainty yields higher expected NPV 's than the deterministic analysis. This is also what Henry (1974) found in his analysis of irreversible decisions under uncertainty.

Table 1: Results of Monte Carlo simulations, $n=1000$.

$\varepsilon_{r0}/\varepsilon_0$	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.9999
$E(NPV)$	4.12	4.39	4.49	4.80	5.55	5.21	5.05	3.81	3.76
$\sigma E(NPV)$	0.22	0.24	0.26	0.27	0.33	0.39	0.40	0.45	0.47
Av. inv. yr. t_d	59	53	48	39	32	26	21	17	13
No. cases	528	547	574	608	653	680	742	807	847

Figure 3 shows expected NPV 's as a function of relative reservation costs. The maximizing value is somewhere around 0.6. As can be seen, the Monte Carlo analysis does not indicate a very precise maximum when the number of simulations is moderate. In the figure, a second order polynomial is fitted to the data points. This is not necessarily the correct function to use. However, the smooth curve is probably a better estimate of expected NPV 's for the interior points than the individual outcomes of the Monte Carlo simulations.

A more intuitive grasp of the results can be obtained from figure 2 where the reservation cost is shown together with five examples of how expected costs develop. If the power plant is started initially, all five outcomes will contribute to an expected NPV of about zero. With the shown reservation costs, the two upper cases with great losses are ruled out. This leads to a higher expected NPV . The cost of this policy is that startup is delayed in the profitable scenarios. In the thick grey line scenario startup is delayed by about 40 years. The solid thick line illustrates a case where an unsuccessful investment would be made if the reservation cost had been somewhat higher. Both in year 5 and in year 35 expected costs for this case are quite low, but in the long run the costs are revealed to be high. This indicates that a decision rule cannot exclude failures completely without ruling out all profitable investments as well. Table 1 illustrates this point. As the initial reservation cost is reduced, the projects are on average delayed further into the future and the number of investment starts is reduced. Eventually expected net present values will start falling when reservation costs are reduced.

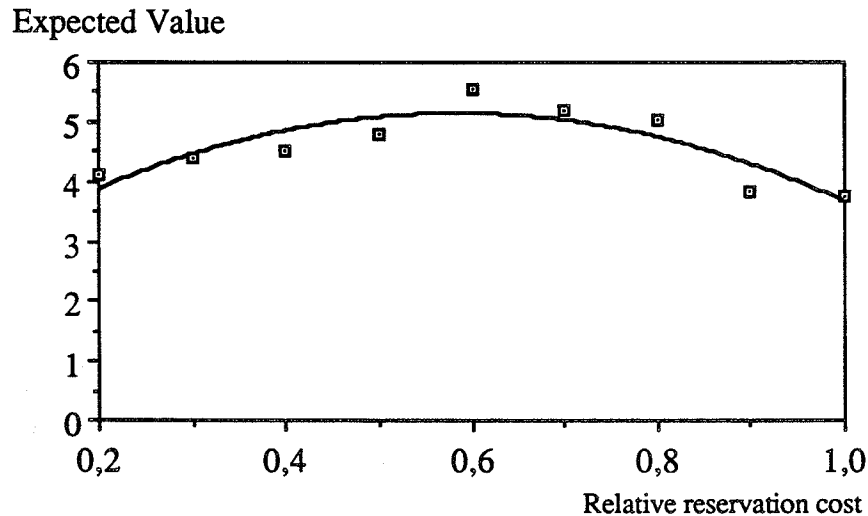


Figure 3: Results of Monte Carlo simulations, indicating the optimal reservation cost.

In the words of deterministic analysis, we have found that the power plant should not be built today unless the project yields a positive *NPV* for "greenhouse" cost estimates that are 40 percent above the expected or most likely costs. If there were no close substitutes to power stations that burn coal, and demand for electric power were inelastic, decisions under uncertainty would have little impact on the actual building of power plants. The reason being that as soon as the supply of power does not grow in pace with demand, prices and benefits will increase to offset the addition to the costs. However, the existence of close substitutes to coal-fired power stations, as well as numerous options for energy saving, implies that a 40 percent addition to the expected "greenhouse" costs ought to have a quite dramatic effect on coal-fired power plant decisions.

Stochastic methods

The Monte Carlo method was the first stochastic method to be applied in the field of economics. Later, dynamic programming and analytical methods have been applied. We comment on these methods in reverse order.

Analytical methods

Stochastic calculus can be used to find analytical solutions to the problem of stochastic optimization. In continuous time, the maximization problem can be stated as:

$$\max_d E_0 \int_{t=0}^T u(x, d) dt$$

The function $u()$ denotes utility, usually including discounting. Decisions or control is denoted by d . The restrictions for the state variable x are expressed by Itô's stochastic differential equation:

$$dx = f(t,x,d)dt + \sigma(t,x,d)dz$$

We notice that the restrictions include an ordinary differential equation f , which contains the mathematics of system dynamics. In addition there is a second term on the right hand side σ which captures the randomness. This term complicates the analytical solution, which typically make use of stochastic integration developed by Itô in 1944 (a generalization of a stochastic integral first introduced by Wiener in 1923), Itô's lemma and Pontryagin's stochastic maximum principle. First in the 1970's the ideas of stochastic calculus have become accessible to the applied researcher through publications like (Åström, 1970; Arnold, 1974; and later Malliaris and Brock, 1987). Early applications in economics are (Merton, 1969; Black and Scholes, 1973).

Analytical solutions to stochastic optimization problems are elegant, exact and provide basic insight into difficult problems. The method is limited to simplified problems. According to (Stensland and Tjøstheim, 1989) much of the work of practical interest use a linear version of the stochastic differential equation, with no influence of time t and control d . "The penalty of this approach is that a fairly narrow class of processes (geometric Brownian motion) emerges that may in fact fail to capture the essential features of the data in a number of cases". In the more general form of the model Stensland and Tjøstheim write that "an explicit solution can only be obtained in certain special cases (Cox and Ross, 1976)".

Dynamic programming

An alternative method is Bellman's dynamic programming approach, which can be illustrated by Søren Kierkegaard's saying that: "Life has to be understood going backwards, and it has to be lived going forwards". Backwards is here from a future date, and living might be interpreted as applying a rule-of-thumb. With dynamic programming the problem is stated in discrete time.

$$V_t(i) = \max_d \left[C(i,d) + \beta \sum_{j=1}^k P_{ij}(d) V_{t+1}(j) \right]$$

Using dynamic programming, one starts solving the problem at the final time T . At this point there is no future value V_{t+1} such that the maximization problem boils down to maximizing the value of the expected immediate return $C(i,d)$ by choosing the best possible decision d . The variable i denotes the state that the earlier development has brought the system into. The maximization has to be made for all possible states i . The resulting expected value for each state $V_T(i)$ becomes the future values $V_{t+1}(j)$ in the earlier period $T-1$. (Future states are denoted by j .) At $T-1$ the expected immediate return $C(i,d)$ is maximized as in the final period. In addition one has to consider the effect of the current policy d on the likelihood of moving from the current state i to a favourable state j at the last point in time. This is modeled by the transition probability $P_{ij}(d)$. Again the maximization has to be made for all possible states i . When the V_{T-1} 's are calculated, one moves to time $T-2$, and so forth to the initial time. Dynamic programming became available in 1957 (Bellman, 1957), and an early stochastic application in economics was by (Samuelson, 1969).

Central to this approach are the transition probabilities P_{ij} , which replace the stochastic differential equation of the continuous approach. The transition probability describes the randomness of the problem. For example, if the system is in the state $i=5$, there is 25 percent probability of getting to a state $j=6$, 50 percent chance of going to $j=5$, and 25 percent probability of ending up with $j=4$. Thus, the transition probability is a matrix. If the process is stationary, one matrix will do, if the probabilities change over time, numerous matrices might be needed.

Using transition probabilities allows for a very detailed description of a random process. However, the size of the problem soon becomes unmanageable with many state variables, non-stationarity, and a fine grid to represent continuous state spaces. According to (Stensland and Tjøstheim, 1989): "A more difficult practical problem is the specification of the transition probability matrix $\{P_{ij}(d), i, j = 1, \dots, k\}$. Usually there will not be enough empirical data so that these quantities can be estimated reasonably well using Markov chain estimates even for moderate values of k ." They proceed to show how a discrete version of Itô's stochastic differential equation or a more general ARMA model can be used to estimate the transition probabilities. This procedure puts severe (and realistic) constraints on the flexibility of transition probabilities, and it allows a limited number of parameters to be estimated from scarce data.

Thus, dynamic programming allows for a more detailed and realistic representation of random processes than the analytical approach. For example, when Stensland and Tjøstheim expands on a simple model, allowing for autocorrelation, the optimal policy is changed considerably (the first example in this paper indicates why).

Monte Carlo Simulations

First note that the Monte Carlo method differs from so-called "scenario analysis"⁹. An early application of the Monte Carlo method for decisions under uncertainty is (Robichek and Horne, 1967). There seems to be few applications of the method lately. The reason is that the method typically relies on heuristics and does not provide optimal solutions. However, the simulation method is still interesting for a number of reasons. The method seems particularly appropriate for researchers with system dynamics background. In the following we discuss how to model an uncertain problem, how to design rules-of-thumb, how to make the rules as efficient as possible, and we comment on the implications for model estimation and implementation.

How to model a problem with uncertainty?

Stochastic optimization involves the modeling of decisions at future points in time when new information is revealed. In the "greenhouse" example a new power plant was built when information turned out to be favourable. Thus, in principle, stochastic optimization involves the modeling of a regular decision process. Typically this process has three steps¹⁰:

Reveal information → Make forecast → Make decision

Information is revealed about random events and state variables, with or without measurement errors and processing delays. A forecast of future development has to be made for the period when the decision has its consequences. Finally, a best possible decision must be based on the forecast. Alternatively the decision might be based on the revealed information directly like in most system dynamics models. In the latter case the last two steps are combined. Below we concentrate on the two first steps, and return to the last step later.

We differ between two cases. First we consider systems with known parameters. These are systems where new information is not expected to improve the parameter estimates. Secondly, we consider problems where measurements of states imply new parameter estimates and better models over time.

An example of a system with known parameters might be a commodity market. Additional time-series data are not likely to give significantly better estimates of price elasticities of supply and demand, effects of inventory coverage on price etc. Randomness is represented in GNP (demand), in weather, political events etc.

If information about price development follows a simple random walk process, the decision process becomes very simple. First, the revealed information comes about by drawing random numbers and adding these to the previous price. Secondly, the forecast is simply equal to the current price, since the random walk process is a martingale¹¹. An example of a decision rule is: Invest if the price is higher than a reservation price.

If information about price development follows a more complicated time-series model like in (Stensland and Tjøstheim 1989) or a system dynamics commodity market model like in (Meadows, 1970), the decision process becomes increasingly complicated. First, randomness is introduced quite easily by drawing random numbers and letting these influence the development of the market model. Secondly, a forecast of market development must be calculated from previous prices in the time-series model, or the different states of the commodity model¹². A decision rule might say: Invest if the expected net present value of a project is greater than zero.

An example of a system with unknown parameters is the "greenhouse" problem. For example, additional time-series data on CO₂ concentration and temperature development will lead to new and improved parameter estimates of the model.

One can, like we have done, simplify matters and assume that information is revealed through a random walk process with declining variances. The choice of process and the exact time development of the variances might be based on simulation experiments with "greenhouse" models and estimation procedures. Martingale properties imply that the last observation is the best forecast, and the decision rule can make use of a time dependent reservation cost.

A more complicated alternative is to include the "greenhouse" model. The parameters of this model must be assumed to be unknown; they are drawn from a priori probability distributions initially. As the simulation gets started, revealed development is used by a built-in Bayesian estimation routine to update the parameter estimates¹³. The initial parameter estimates of this routine equal initial expectations. The updated parameters are used in an equilibrium version of the model to forecast the true costs of the "greenhouse" effect. These forecasts are used in the decision heuristics.

How to construct rules-of-thumb?

"Rules-of-thumb are among the more efficient pieces of optimal decision making" (Baumol and Quant 1964). If a problem is properly stated, nothing is of course more comforting than the optimal solution. However, in cases where it is difficult to state the problem in a simple, solvable form, rules-of-thumb can be fruitful. We give two examples of how rules-of-thumb can be constructed to improve the "optimal" solutions.

First, if an optimal solution has been found to a simplified problem of decisions under uncertainty, this solution provides a good starting point for a rule-of-thumb in a more complex model. The new rule should contain the "optimal" policy as a special case. Monte Carlo simulations will then show whether the "optimal" policy can be improved or if it still holds in the more complicated model. The resulting decision rule will be better than the "optimal" rule. However, we will not know if it is the optimal rule.

Secondly, a rule-of-thumb might be designed to be robust against mis-specifications of the probability distributions of the random variables. According to (Hey 1981, 64): "The use of the word 'optimal' is misleading in that it suggests an objectivity which cannot be present given that the optimality is simply relative to the searcher's prior beliefs. Under very imprecise prior knowledge, reasonable rules may well be better....The general problem of assessing 'how good' are reasonable rules is very similar to the general problem of assessing 'how good' are (reasonable) estimators

....the interested reader could refer to ... (Huber 1972)." See also Bayesian and system dynamics literature on this point¹⁴. Hey's concern about stochastic optimization is similar to Naill's concern about deterministic optimization (Naill 1974).

Searching for the best rule-of-thumb

The statistical nature of Monte Carlo simulations implies that the surface of the expected criterion changes each time the expected criterion is evaluated, see figure 3. This property reduces the efficiency of formal search procedures. Enumeration seems to be a better choice.

Computation time depends on the number of runs in each Monte Carlo simulation and the size of the problem. The number of runs depends on the total variation of the criterion caused by the random variables. The size of the problem can be approximated by the formula $N_{rule}^{S_{rule}}$, where N_{rule} is the number of grid points for each parameter in the rule-of-thumb. S_{rule} is the number of parameters in the rule. In figure 3 there are 9 grid points and one parameter, which implies 9 Monte Carlo simulations. For most purposes 5 grid points would probably be sufficient. To save computation time, the interesting region could be established by using a rough grid and a reduced number of runs in each Monte Carlo simulation. A second iteration could use a fine grid over a small region and a higher number of Monte Carlo simulations. For complex problems computation time might get very long¹⁵. However, Monte Carlo simulations should be more efficient than dynamic programming¹⁶.

Implications for model estimation

Decisions under uncertainty are sensitive to autocorrelation and to the amount of uncertainty in a model. This is different from deterministic decision problems which depend on expected development. Different purposes have implications for the estimation procedures to be used.

Most problems of decisions under uncertainty have been analyzed with very simple processes like the random walk process. According to (Taylor and Kingsman 1978): "It has now been widely accepted that the random walk model describes the stochastic process generating such speculative [commodity market] time series." Taylor and Kingsman go on to discuss four methodological traps that can lead to incorrect acceptance of the random walk hypothesis.

(Stensland and Tjøstheim 1989) have estimated an ARMA model for a problem of decisions under uncertainty. They find the price series to be autocorrelated. However, knowing the long periods of many economic cycles, one should not in general expect that time-series models are likely to reveal autocorrelation from relatively short time-series. Non-linearities¹⁷ and non-stationarity add to this problem.

Scarce time-series data can be complemented by apriori data about a market. This is the approach of Bayesian statistics and in particular of system dynamics, where the required real life interpretation of all parameters facilitates the use of apriori information. (Meadows 1970) has demonstrated convincingly that apriori data like maturity times of hogs, chicken and cattle can be used to explain historical autocorrelation in the prices of the three types of meat.

Estimation techniques that minimize the sum of square prediction errors can lead to misleading results for decisions under uncertainty. This criterion leads to parameters that minimize the variances in one period predictions, which is desired in deterministic analysis. In case of cyclical tendencies in the data and randomness that makes even the true model miss turning-points, square errors might become very high. Minimizing square errors, estimated parameters will attenuate the fluctuating tendency. Fluctuations which might be very important for correct decisions under uncertainty.

Therefore ordinary estimation techniques must be supplemented by other tests, for example checking that the model produces correct amplitudes¹⁸.

Implementing the results of the analysis

In some cases decisions under uncertainty add little to the deterministic analysis. In other cases like in our "greenhouse" example, uncertainty is of great importance. Thus, in general one cannot say that decisions under uncertainty only yield marginal improvements to deterministic decisions. However, the complexity that is added to the analysis implies that one might like such a general statement to be true.

When implementing the results, all the intricate steps of the analysis, whether one uses analytical methods, dynamic programming, or Monte Carlo simulations, do not have to be presented to the client. What the client needs to know is that the assumptions of the model are representative of reality, and that the solution has an intuitive explanation. The first task is similar to that of a regular simulation model. The second task might be eased if the suggested rule-of-thumb takes current practise as a starting point, and shows how current rules can be improved. It might be the case that practical rules already have some provisions for uncertainty.

Finally, when suggesting a decision under uncertainty, one should anticipate that critics of the decision maker might mistake bad luck for bad decision making. One result of this might be that the decision maker denies that bad outcomes are bad. Another result could be that she or he returns to decisions that are in line with what the critics expect. A better strategy would be to announce the next step before uncertainty is revealed. For example, in the "greenhouse" example, a decision to delay the construction of a coal-fired power plant could be accompanied by a statement saying that if costs are revealed to be low, the plant will be built later.

Notes:

¹Thanks to my colleges at the Centre for Petroleum Economics at the Chr.Michelsen Institute and Dag Tjøstheim at the University of Bergen for stimulating discussions and helpful comments.

²The way we use Monte Carlo simulations here differs from the way it is used in "scenario analysis", see note 9. Consider the relationship $y=f(x)$. The probability distribution of x is known, and we want to find the distribution or just the mean value and the standard deviation of y . If the function $f(x)$ is complicated enough an analytical solution might not exist. However, Monte Carlo simulations can always be used. Draw a random number from the distribution of x and calculate y . Repeat this procedure a large number of times. Organize the resulting values of y to construct a histogram of the distribution of y or simply calculate mean and standard deviation from the values.

³The residency time of CO₂ in the atmosphere is thought to be about 100 years, which for practical purposes means that CO₂ releases are irreversible.

⁴A similar assumption about a narrowing distribution has been made by Roberts and Weitzman(1981) in a study of R&D and exploration projects.

⁵Since the random process is discrete in time it can be written:

$$\epsilon_{150} = \epsilon_{145} + v_{145} = \epsilon_0 + \sum_{i=0}^{145} v_i$$

The expected value of ϵ_{150} equals ϵ_0 since the expected values of all v_i 's equal 0. The variance of ϵ_{150} is given as the sum of the variances of the v_i 's. The variance of v_i is $\sigma_{v_i}^2 = \frac{2}{3}(m_0(1-\frac{i}{30}))^3$, where m_0 is half the range of the initial uniform distribution (0.5). Initially $\sigma_{v_0}^2 = 0.083$, and $\sigma_{v_0} = 0.29$. The sum of the 30 variances is $\sigma_{\epsilon_{150}}^2 = 0.67$ and $\sigma_{\epsilon_{150}} = 0.82$. Thus, the final standard deviation is 2.83 times the standard deviation of v_0 . Notice also that while $v(t)$ is uniformly distributed, $\epsilon(t=150)$ is approximately normally distributed because it is a sum of independent random variables.

⁶The delay represents time for the CO₂ releases from the power plant to accumulate in the atmosphere, and it represents time needed to increase the temperature of air, oceans and the ground.

⁷This estimate of the discount rate represents the lower end of estimates based on historical data, see Lind(1982). It represents a high estimate in light of desired future sustainable development, see Moxnes(1989).

⁸In the net present value formula, only 5 percent weight is placed on what happens after year 150.

⁹Using "scenario analysis" one finds optimal policies after each run of a Monte Carlo simulation has revealed all uncertainty. Thus, the solution builds on information that is not available at the time of the decision. Uncertainty only influences the final choice of policy in that the final policy is a blend of the "optimal" policies. According to (Wets 1988, 3): "the major objection to the 'scenario analysis' approach remains the lack of a solid and reliable mathematical basis for the justification of the solutions derived in this fashion." (Rockafellar and Wets, 1987) propose a method that makes "scenario analysis" converge towards the optimal decision under uncertainty.

¹⁰This corresponds to what (Robichek and Horne 1967) does. In an appendix they use explicitly the term "forecast".

¹¹The martingale property is stated as: $E[X_{n+1} | X_1, \dots, X_n] = X_n$. That is, expectations about the future are not influenced by the previous history.

¹²(Forrester 1961, Appendix L) shows how a forecasting procedure based on trend extrapolation can be built into a system dynamics model (and he points to some of the dangers involved when using forecasts). Trend extrapolation is an appropriate procedure to use, if this is in fact the technique that will be used when the actual future decisions have to be made. The parameters of the rule-of-thumb that come out in the end, will reflect the choice of forecasting procedure.

¹³(Sterman et al. 1988) give an example of how a built-in estimation might be modeled in a system dynamics model.

¹⁴See (Zellner 1981; Forrester and Senge 1980)

¹⁵In the "greenhouse" example, each Monte Carlo simulation with 1000 runs took about 20 minutes of computer time on a Macintosh II, using an Excel macro.

¹⁶The complexity of a stochastic dynamic programming problem can be approximated by a similar, but squared formula $N_{states}^2 \cdot S_{variables}$, where N_{states} is the number of grid points for the discretized state variables. $S_{variables}$ is the number of state variables. Typically, N_{states} will be larger than N_{rule} , because state variables that are continuous in reality cannot be forced to jump in large steps. This would change the nature of the random process and the optimal solution. S_{rule} should probably be as close to $S_{variables}$ as possible. Thus, the facts that the dynamic programming formula is squared, and that N_{states} typically is larger than N_{rule} , imply that the Monte Carlo method will be more efficient than dynamic programming when the number of state variables gets large.

¹⁷For a discussion of non-linear models see (Tjøstheim 1986).

¹⁸See note 14 for references to such techniques.

References

- Arnold, L., 1974: *Stochastic Differential Equations: Theory and Applications*. John Wiley & Sons, New York.
- Baumol, W.J. and R.E.Quandt, 1964: Rules of thumb and optimally imperfect decisions. *American Economic Review*, 54, 23-46.
- Bellman, R., 1957: *Dynamic Programming*. Princeton University Press, Princeton, New Jersey.
- Black F. and M.Scholes, 1973: The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637-54.
- Cox, J. and Ross, S., 1976: The valuation of options for alternative stochastic processes. *Journal of Financial Economics* 3, 145-166.
- Forrester, J.W. and P.M. Senge, 1980: Tests for Building Confidence in System Dynamics Models, in *System Dynamics*, TIMS series on Studies in Management Science, American Elsevier Publ. Co. NY.
- Forrester, J.W., 1961: *Industrial Dynamics*. MIT Press, Mass. USA.
- Henry, C., 1974: Investment Decisions Under Uncertainty: The "Irreversibility Effect". *The American Economic Review*. Vol.64, No.6, 1006-1012.
- Hey, J.D., 1981: Are optimal search rules reasonable? And vice versa?(And does it matter anyway?). *Journal of Economic Behaviour and Organization* 2, 47-70.
- Huber, P.J., 1972: Robust statistics, A review. *Annals of Mathematical Statistics* 43, 1041-1067.
- Lind, R.C. et al., 1982: *Discounting for Time and Risk in Energy Policy*, Resources for the Future, distributed by The Johns Hopkins University Press, Baltimore and London.
- Malliari, A.G. and W.A.Brock, 1987: *Stochastic Methods in Economics and Finance*. North Holland, Amsterdam.
- Meadows, D. L., 1970, *Dynamics of commodity production cycles*, Wright- Allen Press, Cambridge, Mass.
- Merton, R.C., 1969: Lifetime portfolio selection under uncertainty: The continuous-time case. *Review of Economics and Statistics* 51, 247-57.
- Moxnes, E., 1989: Sustainable Development. Chr. Michelsen Institute, N-5036 Fantoft, Norway, CMI-no. 30160-1.
- Naill, R.F., 1974: *Optimizing Models of Social Systems*. DSD-15, Thayer School of Engineering, Dartmouth College, N.H. USA.
- Roberts K. and M.L.Weitzman, 1981: Funding criteria for research, development, and exploration projects. *Econometrica*. Vol.49, No.5, 1261-1288.
- Robichek, A.A. and J.C.vanHorne, 1967: Abandonment value and capital budgeting. *The Journal of Finance*. Vol.22, No.4, 577-589.
- Rockafellar, R.T. and R.J-B Wets, 1987: Scenarios and policy aggregation in optimization under uncertainty. WP-87-119, IIASA.
- Samuelson, P., 1969: Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics*, 51, 239-246.
- Stensland, G. and D.Tjøstheim, 1989: Optimal Investments Using Empirical Dynamic Programming with Application to Natural Resources. *Journal of Business*. vol. 62, no.1, 99-120.
- Sterman, J.D., G.P.Richardson, and P.Davidsen, 1988: Modeling the Estimation of Petroleum Resources in the United States. *Technological Forecasting and Social Change* 33, 219-249.
- Taylor, S.J. and B.G.Kingsman, 1978: Non-stationarity in Sugar Prices. *Journal of the Operational Research Society*. Vol.29, 10, 971-980.
- Tjøstheim, D., 1986: Estimation in nonlinear time series models. *Stochastic Processes and Their Applications* 21, 251-273.
- Wets, R.J-B, 1988: The Aggregation Principle in Scenario Analysis and Stochastic Optimization. Working Paper, Department of Mathematics, Univ. of California, Davis.
- Zellner, A., 1981: Philosophy and Objectives of Econometrics, in Currie, D., R.Nobay and D.Peel (eds.): *Macroeconomic Analysis: Essays in Macroeconomics and econometrics*, Croom Helm, London.
- Åström, K.J., 1970: *Introduction to Stochastic Control Theory*. Academic Press, Inc., New York.