RULES OF THUMB AND REAL OPTION DECISION BIASES FOR OPTIMALLY IMPERFECT DECISIONS: A SIMULATION-BASED EXPLORATION

Burger-Helmchen Thierry*

Abstracts

Investment decisions about an uncertain project are a difficult task. A decision maker can use calculation techniques such as net present value or real option. The accuracy of the technique employed can provide a significant modification to the final decision. Each of these techniques makes the assumption that the decision maker acts in a neutral way without any cognitive bias, such as overconfidence in his or her opinion. Much research in behavioural finance show that pessimistic or optimistic feelings of the decision maker can potentially lead to wrong decisions. We explore the relation of some decision biases and the use of evaluation techniques. By employing simulations, we show that the choices of a specific technique can emphasise or reduce decision bias and investment errors.

Key words: decision-making, real options, uncertainty.
JEL Classification: D81, G3, M20.

1. Introduction

"In planning major initiatives, executives routinely exaggerate the benefits and discount the costs, setting themselves up to failure."

D. Lovallo, D. Kahneman, 2003.

"Society values risk taking, but not gambling, and what is meant by gambling is risk taking that turns out badly."

March J.G., Shapira Z., 1987.

Kahneman and Lovallo pinpoint that decision makers routinely make errors in their decisions. This paper investigates how rules of thumb that are routinely utilised by decision makers can influence these errors. We compare different types of rules of thumb, including real options and a deflating criterion to show which kind of rules of thumb leads to exaggeration in profit estimation.

On the one hand, it is easy to jump to the conclusion that the widespread use of rules of thumb is good evidence of the sloppy workmanship on the part of business management. On the other hand, if the rules of thumb are historically in use and neutral, they can eliminate some decision bias. In this paper, we argue that rules of thumb can be among the most accurate tools of decision-making in some cases and in others lead to a critical decision bias.

A simulation procedure is used to determine the relevant properties of a number of alternative pricing rules of thumb and evaluation techniques to compare their performance to help decision-makers in a monopolistic situation.

The paper is organized as follows: First, we discuss some biases that can occur when a decision maker evaluates a project. Second, we describe our simulation model and explain how we incorporate real options reasoning and decision correction into it. We then provide the results of our simulation. Finally, we discuss the implications of our study.

2. Optimally Imperfect Decisions and Decision Biases

For the rest of this paper we will use as an example a project of technology producing a new product in a monopoly situation and the choice to invest in it or not. The errors that can occur

* BETA, University Louis Pasteur, France.

in such a choice are twofold. The decision maker can accept a project that turns out to be a bad one, which means the project's ex post is not profitable and should have been rejected. Or, a second error includes the decision makers rejecting a project that should have been accepted.

The standard case in corporate finance textbooks for deciding to invest or not in a project presents a decision maker with perfect information. The trick for him or her is to perform a calculation and decide to choose the project with the highest profit, which is usually calculated with Net Present Value (NPV). When we introduce risk or even uncertainty in the form of imperfect information or inexistent information, the calculation becomes more difficult and the decision process is less straightforward. The more refined the decision making process becomes, the more expensive it is likely to be, notably with the introduction of new information search processes. This later example is similar to the real world where it is impossible to perform any perfect estimation of an investment opportunity. Whatever the techniques of evaluation a decision maker uses, a calculation down to the last decimal place is pointless in any event.

Accepting that it is impossible to collect and compute all information necessary to obtain a perfect decision, we can define criteria for a decision to be an optimally imperfect decision: the marginal cost of additional information gathering or more refined calculation must be equal to its marginal gain in revenue. Unfortunately, this definition is not very useful in practice and there must be other "limits" to the search of new information. Simon's (1979, 1991) opinion with reference to this question is that people usually act in a way to satisfy themselves with the situation and not to maximise it. The satisfaction of a decision maker has an idiosyncratic concept and opens the door to numerous biases.

The importance of analysing the effects of judgmental processes and cognitive biases on strategic management has been examined by several authors (Hoskisson et al., 1991; Schwenk, 1984) including the entrepreneurial specific biases and the role real options can play (Mcgrath, 1999). While part of the effects stems from the underlying structure and features of the human judgmental system (Tversky and Kahneman, 1974), other potential determinants of "biased" decisions are the incentives and penalties associated with the outcomes of decisions. An executive's decision in the situation of uncertainty appears to be affected more by attempts to avoid failure rather than by attempts to pursue risky alternatives whose potential success can be described only as probabilistic (Shapira, 1994). There is a tendency for decision makers to behave in a risk adverse manner. Pursuing this avenue is manifested by attempts to continue building on alternatives that produce small but certain profits rather than pursue new opportunities. As March (1991) noted, firms often get caught in a conflict between the need to explore new opportunities and the tendency to exploit existing resources.

To choose between projects several considerations are important. Decision makers need an appropriate measure to discriminate between projects that should be accepted from those that should not be rejected. Likewise, decision-makers need an appropriate measure to discriminate between success and failure. While it may be easier to employ a common measure across time, the realities of the technology development process, communication to shareholders, etc., compel decision makers to choose measures that are different across time.

For example, net present value (NPV) and internal rate of return (IRR) are well-accepted selection measures to evaluate projects (Brealey and Myers, 2002; Luenberger, (1997). In contrast, return on investment (ROI), profitability, and market share are widely used measures of success. In addition to choosing measures, decision makers need to choose an appropriate criterion for each measure to discriminate between the accepted and rejected regions as well as between success and failure. Since decision makers seek to increase their ability to choose technologies that are likely to be successful, they also need to establish the relationship between selection criteria and indicators of success.

Decision makers and project initiators may represent the value of a project differently (Williamson, 1975). For instance, project initiators may be overly optimistic and aggressive in promoting the benefits of their project and forecasting high demand, while being blind or silent about the potential costs involved. Differentiating the standard deviation of the project in two different measures, one for the demand and the other for the cost, is a first step toward differentiating different biases (overconfidence, and neglecting).

3. Optimal Rules of Thumb: Definitional and Calculation Problems

3.1. Choosing a Rule of Thumb

Before doing any calculation, the decision maker has to choose between different rules of thumb. This choice is not an easy task and introduces some biases. Conceptually, the ranking of rules of thumb is not easy because problems of multidimensionality are inherently involved. At first glance it may seem sufficient to say that between two rules that are equally costly to operate, the one that yields the closer average approximation to the true maximum is preferred. But some measure of this dispersion of the results is, surely, also of comparable relevance. In fact, all the characteristics of the frequency distribution of errors are significant. The "best" rule of thumb can only be determined by taking all these characteristics into account and assigning appropriate weights to them. It appears that there is no cut and dry mechanical way to determine which of two rules is the more satisfactory unless one of them is what might be called "Pareto preferable"; that is, it is not inferior in any of its relevant characteristics and superior at least in some.

The decision bias that emerged here can be either voluntary or involuntary. Laverty (1996) provides a description of some biases that can emerge in that phase, which insist mostly on the short-term. The bias in the choice can be managerial opportunism, if one takes a rule of thumb that favours a project that has the preference of the decision maker (here the manager), financial market short-term who obligates the decision maker to invest in a project with early cash flows, or simply the incompetence of the decision maker. Note that Laverty pinpoints the importance of the environment; in this work we partly neglect the environment in choosing a non-contestable monopole situation.

3.2. Determining an Optimal Rule of Thumb

We follow Baumol and Quandt (1964) who pioneered the use of simulation techniques for testing rules of thumb. We justified our approach and use of simulations in the same way. One of the crucial steps in the determination of an optimal rule of thumb is the calculation of the expected consequences. Ideally, this step should produce a frequency function of the errors to be expected of any such rule of thumb in its approximation to true maximal solutions. In principle, it may be possible to analytically determine these frequency functions if the structure of the system and the probability distributions of the structural variables are known. However, in practice such a procedure is almost impossible. First, a priori specification for the structural variables usually represents most heroic assumption that rarely stands up under comparison with the facts. Secondly, even under these conditions, analytic methods of determination of the expected effects of a rule of thumb are likely to be prohibitively difficult. This means that one must turn to methods of testing and calculation. The methods of simulation using artificial data seem well suited to the problem. By considering a sufficient high number of simulated cases, the analyst may be able to come away with a fair idea of the rule's likely performance.

An attempt to use simulation and test the accuracy of real option against standard decision criteria's is made by Kumaraswamy et al. (2003). They explore whether the use of evaluation and organizational routines that help a firm to recognize and realize the value of flexibility options attached to its R&D opportunities/projects pay off in terms of improved R&D performance. Specifically, they use a computer-simulation-based approach to compare the baseline (or traditional) version of two evaluation techniques – the Net Present Value (NPV) technique used in finance and the Multifactor Evaluation Process (MFEP) technique used in decision analysis – with their respective options-enhanced versions. Their simulation results indicate that the use of the option-enhanced version of either technique resulted in a significantly higher success rate of funded projects and overall R&D productivity when compared to that achieved using the baseline version. An implication is that the consistent application of straightforward evaluation and organizational routines can capture the intuition behind complicated real options mathematics and yield improved R&D performance.

4. Designing a Price-Setting Rule of Thumb and an Evaluation Rule

Two levels of rules of thumb are used. Firstly, the decision maker must decide the price and adequate quantity for a new product. Ignoring the exact demand and cost function, the decision maker first has to estimate this demand and cost function. Secondly, he has to estimate the expected profit and decide if he will continue with this project or prefer a risk-less investment with a small profit.

The detailed discussion of our basic experiments can be conveniently divided into three parts: the methods that were used to generate demand and cost functions (3.1), the rules of thumb that were employed (3.2), and the method that was used to obtain estimates of the performances of various rules of thumb (4).

4.1. Description of the Experiments: the Cost and Demand Function

The demand (average revenue) curve facing the decision-maker and his total cost curves are generated by a random process. The resulting pair of demand and cost curves is regarded as the "true" curves, correctly describing the actual state of affairs. The decision-maker is assumed to know only two points of these curves. He uses some rules of thumb to calculate from these two pairs of points the price he is to charge in the market. We can determine from the true demand and cost curves the profit maximizing price and the maximum amount of profit as well as the decision-makers actual profit resulting from his rule of thumb price. The comparison of actual and maximum profit provides an estimate of the efficiency of the rule of thumb.

4.1.1. The Demand (Average Revenue) Function

The hypothetical product we deal with is assumed to be such that the quantity demanded is zero when the price, p, is greater or equal to 21. The demand function (as well as the cost function) is discrete. The quantity demanded is defined only for integral values of p. The demand function was obtained by requiring the graph to start off at the point with coordinates p = 21, q = 0 and the calculating increments in the quantity demanded corresponding to successive unit reduction in the price. The increment was selected in one of two ways[1]: (i) in the runs 1, 2, 3, 4 and 5 the increments are uniformly distributed over the integers 1,.., 64; (ii) in Run 6 the tth increment is Xt + 2t where Xt is uniformly distributed over the integers 1,.., 64. Note that the total revenue function needs not be concave everywhere, and that a finite amount is demanded when the price is zero.

4.1.2. The Total Cost Function

The demand function was obtained by associating with each price, p, a quantity q = f(y). The cost function associates with each of these q's a total cost figure C = g(q). With zero output we associate a cost figure that is chosen from the uniform distribution over the integers 1, …, 64. Each successive cost figure (associates with a successively higher output level) is obtained from the immediately preceding one by the addition of an increment. Depending on the run, the tth increment is (i) chosen from the same distribution, (ii) a multiple of an increment chosen from the same distribution, or (iii) expressible as Yt + 8t where Yt is chosen from the same distribution.

It is obvious that maximum profit and the optimal price are easily determined from these functions simply by examining each of the 22 possible price levels (p = 0 to p = 21) and calculating, in each case, the associate profit figure. The six experiments are summarized in the table below, showing the manner in which the tth increments for the demand and cost functions were generated. Letting Xt and Yt be random variables distributed uniformly over the integers 1,…,64, we have the following:

[1] These two methods generate demand function whose expected values are, respectively, linear and quadratic. That is, in Runs 1-5, with each unit increase in price the expected (average) decrease in demand is (1+2+…+64)/64= 32,5. Hence, the expected *dq/dp* is constant and the expected demand curve is linear. In Run 6, the expected sales increment when the price is p = 21-t is Δq = 32,5 + 2t, so that we can take $\Delta q/\Delta p$ = -32,5-2t (approx.), and hence the expected demand function generated in Run 6 is quasiquadratic. Similar interpretations hold for the cost functions that were used in the various runs.

Table 1

Demand and cost function increment

		Demand function increment	Cost function increment
Run	1	Xt	Yt
	2	Xt	4Yt
	3	Xt	8Yt
	4	Xt	Yt+ 8t
	5	Xt	4Yt+ 8t
	6	Xt + 2t	Yt

We considered four possible types of approximating functions: linear, quadratic, logarithmic, and exponential. Of course various combinations are possible; the demand curve can be linear, etc. Thus, 16 combinations are possible. In the following we limit to only four of these combinations. We feel that four combinations, with the different options simulated and leading to 60 cases to compare are sufficient for this exploratory work. We also include two other very naïve rules as standards of minimal performance. The different rules for defining the price p are the following; the detailed calculation that rules imply are given in the appendix.

Table 2

Price in demand and cost function

		Demand function	Cost function
Rule	1	P is fixed to 11 without any estimation	
	2	P is chosen randomly	
	3	Linear demand p = a – bq	Linear cost c = d + eq
	4	Linear demand p = a – bq	Quadratic cost c = dq+eq2
	5	Linear demand p = a – bq	Logarithmic cost c = dq + e log (q+1)
	6	Hyperbolic demand p = a-b/q	Quadratic cost c = dq+eq2

We imposed one rather strong condition on the construction of our rules in order to assure that they were sufficiently simple to qualify as rules of thumb: in no case was the decision maker asked to take account of more than two points on his demand curve and two points on his cost curve in his calculation. This can be seen as a very strong limitation. However, we think that major prospective studies that a decision maker can have give a “level of expected demand or cost” and do not give a perfect shape of the curve. In that respect we think that these two point’s approximation is a rather realistic one.

4.2. Description of the Experiments: the Profit Evaluation Rules

The original work of Baumol and Quandt (1964) use a single period estimation of the profit. We notice that major bias in decision occurs in the way decision makers compute the profit, not only in the initial period but also the evolution of the profits in the following periods. It is difficult for a decision maker to be convinced that the profits would be enormous in the first period, thus they are more likely to believe that the profit is modest in the first period and will increase in time. We chose the simplest multi period possible: two periods. Two periods are necessary for the use of the options that we describe. We introduce new information in the following manner. The real demand and cost functions do not change, and the decision makers obtain two new points for performing their calculation. The points given to the decision maker have the property to be more separate from one another than the points in the first period. This will lead to a more exact estimation of the maximum profit in the majority of simulations. It is important to note that this is not always the case; we ensure that decision makers can make the two types of errors in calculating the profit in the second period.

Our proposition is to investigate the effects of including real option as a rule of thumb. If the results obtained by using real option will lower the investment errors made by decision makers (counter bias) or in opposite, the real options do increase the decision bias. An example of this logic can be found in reading real option standard presentations (Trigeorgis 1996; McDonald, 2000).

The overwhelming majority of examples are built in the same manner. An investment project is presented, but unfortunately the NPV of this project is negative and decision makers should not invest. The real option rhetoric consists in introducing flexibility in the future employment of the project. The capacity a decision maker has to "manipulate" an investment during its lifetime is to take advantage of good situations and avoid bad situations when new information is obtained. Of course, this option has a value that in addition to the previous NPV will give a positive expected value, and so categorize the investment into a profit to be undertaken as an investment.

This example, widely used, builds on the capacity of decision makers to adequately forecast the future options. We can also see in this description the following procedure if one cannot invest into a project based on standard NPV criterion, add some "options" and obtain a profitable project. This represents a standard moral hazard opportunism case where decision makers can make a project more attractive. It is this perspective using real option that magnifies decision-makers bias.

The growth option can be an expression of bias, like overconfidence in the future. The option to wait can show reluctance to invest or fear of doing a mistake. This fear can lead to underinvestment and the loss of profitable projects.

Notice that the same option can be used in a contrary example to reduce the same bias. Taking the growth option into account can annihilate the bias of fear, and in the same manner including options to wait in the project can be a good correction to the enthusiastic project leader calculation.

We compute the growth and wait options in the following manner:

The growth option. Once the project is undertaken and the decision-maker obtains the first period profit, he can legitimately await a growth in the profit for the second period. We adopt the following decision rule for the decision maker: if the expected profit in the first period is lower than the minimal profit when he invests in a risk-less project, then he expects a growth in future projects, and the price for the second period is based on the second period estimation.

The option to wait. Here we again put into balance a minimal return from a risk-less project and the expected profit of the project in evaluation. The decision rule is the following: if the expected profit in the first period is lower than the risk-less project, then investing in risk-less project for the two periods (wait for the risky project) takes place. If the calculation of expected profits in the second period is above risk-less projects profit, investing in the project occurs.

The KTL rule. We used a third rule based on work by Lovallo and Kahneman (2003), Kahneman and Tversky (1979). These authors present a method for evaluating investments using a procedure introducing an outside view. This procedure was built to adjust the optimism of decision-makers by lowering the expected profit. The procedure uses three variables. The first is the expectation of the profit made by the decision-maker; E (πi), we use the expectation of profit in the period i. The second variable is the use of a generally admitted profit for an investment project of the same reference class. Here we use the same risk-less project for the option computation. One can argue that this project is risk-less, so the class is not the same. Here, we interpret the class of projects in a more general way, thus all the projects other than the decision maker can be undertaken in the same periods than the risky project under evaluation. We use the letter r to describe this risk-less project profit. The third variable is the correlation between the two projects ρ.

Kahneman-Tversky and Lovallo propose the following formula to correct the profit expectation:

$$E(\pi_1) + \left[\rho\left(r - E(\pi_1)\right)\right] = KTL \tag{1}$$

We design by KTL the value obtained. If the sum of KTL for the two periods is lower than the obtaining profit r in the two periods, decision makers decide to invest in the risk-less project.

5. Methods of Testing

We use two rules (1 and 2) to estimate the price yield results that are independent of the particular information assumed to be immediately available to the decision maker. Therefore, only one rule of thumb solution is provided by each of these two rules for any randomly given generated demand function-cost function pair.

This is not the case with Rules 3, 4, 5, and 6. The particular solution obtained from these rules depends on which two out of the possible 22 points on the demand and cost functions are used in fitting the demand and cost functions.

We have assumed that a zero output or an output saleable on a zero price is never used by the decision maker in this calculation. Hence, 20 points remain out of which "nature" was taken to choose two for the decision maker to observe. In order to evaluate the performance of a rule of thumb, we calculated its solution for each possible pair of points chosen out of the total 20. There are 190 such pairs; therefore, there are 190 profit figures corresponding to a single "true" demand function-cost function combination per period.

For each particular true demand function-cost function combination, we obtained the profit accruing under Rules 1 and 2, and each of these was compared with the average profit figure accruing under Rules 3, 4, 5 and 6 and to the different rules to estimate the profit. Further details are given in the appendix.

Baumol and Quandt conducted 24 simulations for each run. Because of a dramatic increase in calculation since the mid sixties, we chose to compute 24.000 simulations for each of the six runs. Tables 1 to 6 give the results of this simulation. In the column, the standard deviation of each mean profit is given in parentheses. We had the growth, wait option and KTL to measure a mean profit realisation that is a simple addition of the profit in the two periods when the price is set according to the information available at this point without any actualisation for the second period (interest rate = 0), which corresponds to the results ex-post. Also, the expected row corresponds to the sum of the expectation in the two period's ex-ante. Numerical details of the simulation runs are given in the appendix.

6. Results of the Calculations

Three thresholds should be kept in mind while reading Tables 1 to 6. First, the minimum and risk-less profit is 2000 per period of simulation, which makes every profit obtained under 4000 for the two periods of simulation to be considered a very bad choice for the decision maker. The second threshold is the naïve Rule 1, who consists in always charging the same price p = 11. Performing less than the profit of Rule 1 is also a sign of poor decision. Notice that Rule 1 in some of our simulations performs very well. This corresponds mainly to the restriction imposed to our model (price limited between 0 and 21). However, Rule number 1 is introduced to help screen the results. The third threshold is the first column of result: the maximum profit one can obtain in investing in the risky project. Note that it is possible to obtain a higher profit in some simulations by using the risky project in one period and deciding to invest in the risk-less project in the other period.

The results in Tables 1 to 6 suggest that Rules 2 and 6 can be eliminated immediately on the basis of their average poor performance in comparison to the other rules. Rule 2 performs well in Run 1, but performs badly in the following Runs. This is due to the upward shift in the cost function in Runs 2-5, thus the average optimal price is higher in these runs.

Rules 3, 4 and 5 give globally comparable results, but with a decrease in the ration mean profit/maximum profit when we look from Rule 3 to Rule 4. This is because the costs are expected to be higher in these rules.

A comparison between the expected profits and real profits show that the decision maker expected more than the real outcomes in the majority of the cases. Only in some runs using Rule 6 the expectations are smaller than the real results. This is mainly because the demand function of Rule 6 is hyperbolic and the price limitation that we employ limits us to a relative flat part of the demand. This high expectation can be seen as the sum of the decision-maker's bias. The difference between expected profit and real profit can turn out to be extremely high in some cases.

The comparison between the growth option case and the mean profit without any option calculation shows us that the decision maker using growth option outperforms the standard case.

But the standard deviation of growth option is high in comparison to the standard case; it is almost double. Also, many results of growth options are under the minimal 400 threshold. What this means is that this rule pushes the decision maker to not use the risk-less investment possibility and turn to the risky project more easily.

Comparing the option to wait case and the procedure proposed by Kahneman, Tversky and Lovallo, we can notice a high similitude between these two groups of results. This leads us to say that using an option to wait is a good way to reduce too enthusiastic expectations from decision-makers. The option to wait is a counter-bias effect. This effect is desirable when the expectations are far too high. Unfortunately, this procedure also seems to favour projects instead of investing in the risk-less projects; thus, the result is that in many cases the option to wait and KTL procedure gives result under the 4000 minimum.

It is also noteworthy that the performance of these rules of thumb are in terms of mean profit globally high if we take into consideration that the profit calculations were based on demand and cost functions fitted only to pairs of points on the two functions.

Notice also that the real profit obtained in the second period is more than 80 percent of the case greater than the profit obtained in the first period. This is because we specify that the points taken for the calculation are more separate than in the first period. Since the demand and cost functions are linear or quadratic, this phenomenon is normal and we could expect better performances when the points are far one from each other.

7. Concluding Remarks

Kogut and Kulatilaka (2003) make us aware that decision bias can come not only from individual decision makers, but also from corporate pressure. They remark that the option models have been moved from financial markets to corporate decision making only by searching to fit different domains together (financial market and corporate decision) without taking into consideration the behavioural decision-making biases introduced. However, they suggest that given the negative evolutionary consequences in ignoring options like investments, organizations invent heuristic rules to counter these biases. They conclude by proposing the idea of a domain translation that shows how the basic insight of option pricing can be preserved through evolving complementary organizational rules, such as increasing the frequency of monitoring the value of an investment. Again, a simple option-pricing simulation illustrates the joint influence of a status quo bias and frequent monitoring.

We separately test real options, making the strong hypothesis that the decision biases are also separable. It is wrong to implicitly assume that behavioural biases do not themselves “duel” with each other, as do options, or that competition between decision makers in a firm will not lead to increase rather than decrease bias, or that organizations are unable to develop capabilities to counteract these pathologies. Folta and O’Brien (2004) examine the tension between the options to defer, and the options to grow and uncertainty. They investigate these relations on the decision established by firms to enter a new industry.

Nevertheless, this work shows that there is a link between real option and decision bias. It encourages further work to describe more accurately the effects of the decision makers behaviour and their choice of rules of thumb including real options.

References

1. Baumol W.J., Quandt R.E., “Rules of Thumb and optimally imperfect decisions”, American Economic Review, 1964, N°4. – pp. 23-46.
2. Brealey R.A., Myers S., Principles of Corporate Finance, McGraw-Hill, 2002.
3. Folta T.B., O’brien J.P., “Entry in the presence of duelling options”, Strategic Management Journal, 2004, N° 25. – pp. 121-138.
4. Hoskisson R.E., Hitt M.A., Hill, C.W.L., “Managerial risk taking in diversified firms: An evolutionary perspective”, Organization Science, 1991, N°3, – pp. 296-314.
5. Kahneman D., Tversky A., “Intuitive Predictions: Biases and Correction Procedures”, Management Science, 1979, N°12. – pp. 313-327.

6. Kogut B., Kulatilaka N., “Real Option Pricing and Organizations: The Contingent Risks of Extended Theoretical Domains”, Academy of Management Review, 2004, No.1. – pp. 102-110.
7. Kumaraswamy A., Redmond M., Baveja A., “Exploring the effect of real options reasoning on R&D, opportunities / Projects evaluation: a simulation study”, 2003, presented at the International Conference on the Management of Research and Development, New Delhi.
8. Laverty K.J., “Economic 'Short-Termism': The Debate, The Unresolved Issues, and the implications for Management Practice and Research”, Academy of Management Review, 1996, No. 3. – pp. 825-860.
9. Lovallo D., Kahneman D., 2003, “Delusions of Success”, Harvard Business Review, July, pp. 56-63.
10. Luenberger D.G., 1997, Investment Science, Oxford University Press.
11. March J.G., 1991, “Explorations and exploitation in organizational learning”, Organization Science, Vol. 2, No. 1. – pp. 71-87.
12. March J.G., Shapira Z., 1987, “Managerial perspectives on risk and risk taking”, Management Science, Vol. 33. – pp. 1404-1418.
13. Mcdonald R., 2000, “Real Options and Rules of Thumb in Capital Budgeting”, in Project Flexibility, Agency and Competition: New Developments in the Theory and Application of Real Options, (eds.) Brennan M., Trigeorgis L., Oxford Press.
14. Mcgrath R.G., 1999, “Falling Forward: Real Options Reasoning and Entrepreneurial Failure”, Academy of Management Review, Vol. 21, No. 1, – pp. 13-30.
15. Schwenk C.R., 1984, “Cognitive simplification processes in strategic decision making”, Strategic Management Journal, Vol. 5, – pp. 111-128.
16. Shapira Z., 1994, Risk Taking: A managerial Perspective, New York, Russell Sage Foundation.
17. Simon H.A., 1979, Models of Thought, Yale University Press.
18. Simon H.A., “Bounded rationality and organizational learning”, Organizational Science, 1991, N° 2. – pp. 125-134.
19. Trigeorgis L., Real Options. Managerial flexibility and Strategy in Resource Allocation, The MIT Press, 1996.
20. Tversky A., Kahneman D., “Judgement under Uncertainty: Heuristics and biases”, Science, 1974, N° 185. – pp. 1124-1131.
21. Williamson O.E., Markets and Hierarchies, New York: The Free Press, 1975.

Appendix 1: Decision-maker estimation of the demand and cost function

For the purpose of calculating rule of thumb solutions to the profit maximisation problem, we assume that the decision-maker knew exactly two points on his demand function and two points (with the same abscissa as the point of the demand function) on his cost function. The decision-maker was then assumed to calculate his pseudo-optimal price on the basis of the information contained in these four points. It is to be noted that if the pseudo optimal price, p, turned out to be fractional, it was rounded to the smallest integer greater than p. The profit accruing to the decision maker under a particular rule of thumb was then calculated from the true demand and cost functions.

It is clear that the decision-maker estimates of the demand and cost functions will depend, in general, upon which two of many possible points on the demand and cost functions the decision maker is assumed to know. In order to evaluate the average or expected effectiveness of a rule of thumb statistically, the experimenter must present the decision maker, seriatim, with various alternative pairs of points and observe his behaviour in each of these possible cases.

The six rules of thumb that were investigated were essentially chosen on the basis of their (relative) simplicity. No attempt was made a priori to choose rules that would in some sense be highly reasonable. In particular, the first two rules that are about to be described may be considered totally arbitrary and unreasonable. These "naïve rules" were designed only to provide a minimum standard of performance in that any rule that does not provide results better than those offered by our naïve rules should, doubtless, be rejected out of hand.

(Naïve) Rule 1: Fixed price. Irrespective of the location of the two known points on the demand function and the cost function, charge a price p = 11, which is the midpoint of the range of relevant positive prices. Strictly speaking, this is not even a rule of thumb under our definition because no decision machinery is provided whereby the decision maker can actually select his unvarying price on the basis of objectively measurable data.

(Naïve) Rule 2: Random price. Choose a price from the uniform distribution over the integers 3, …,17, irrespective of the two known points from the demand and cost functions.

Rule 3: Linear demand, linear cost. Fit to the two known points on the demand function the linear relation p = a – bq, and to the two points on the demand function the linear relation c = d + eq. If the two pairs of points are given by (p1, q1), (p2, q2) and by (c1, q1), (c2, q2), estimates of points a, b, d, and e are obtained from the expressions:

$$\hat{b} = \frac{p_1 - p_2}{q_1 - q_2}\,, \tag{2}$$

$$\hat{a} = p_1 + \hat{b}q_1\,, \tag{3}$$

$$\hat{e} = \frac{c_1 - c_2}{q_1 - q_2}\,, \tag{4}$$

$$\hat{d} = c_1 + \hat{e}q_1\,. \tag{5}$$

For example, here the expressions for b and â are obtained directly by elimination from the two equations p1 = â – bq1 and p2 = â – bq2. The decision-maker is then assumed to calculate his pseudo-optimal price by maximising profit on the basis of the estimated demand and cost function. Profit is:

$$\pi = aq - bq^2 - d - eq\,, \tag{6}$$

and maximising

$$\frac{d\pi}{dq} = a - 2bq - e = 0 \tag{7}$$

so that:

$$q = \frac{a-e}{2b} \quad (8)$$

and, from our demand equation,

$$p = a - b\left(\frac{a-e}{2b}\right) = \frac{a+e}{2} \quad (9)$$

By virtue of the method by which the true demand function is generated, b > 0 and the second order condition for a maximum is always satisfied.

Rule 4: Linear demand, quadratic cost. In Rule 4 the demand function was estimated in the same manner as in Rule 3. The cost relationship fitted was the quadratic function c = dq+eq2. The coefficient estimates were given by:

$$\hat{e} = \frac{c_1 q_2 - c_2 q_1}{q_1 q_2 (q_1 - q_2)} \quad (10)$$

$$\hat{d} = \frac{1}{q_1}\left(c_1 - \hat{e} q_1^2\right) \quad (11)$$

$$\pi = aq - bq^2 - dq - eq^2 \quad (12)$$

yields the first order condition

$$\frac{d\pi}{dq} = a - 2bq - d - eq = 0 \quad (13)$$

whereby

$$q = \frac{a-d}{2(b+e)}, \qquad p = a - b\left(\frac{a-d}{2(b+e)}\right) \quad (14)$$

The second order condition is:

$$2(b+e) > 0 \quad (15)$$

We observed the following conventions:

1) Where the second-order condition failed (which can happen since e may be negative), a price p = 21 was charged[1].

2) If the calculated optimal quantity was non-positive, a price p =21 was charged.

3) Marginal cost was constrained to be positive for the entire relevant range of outputs.

Rule 5: Linear demand, logarithmic cost. In Rule 5 the demand curve was obtained as before. But this time the cost function was c = dq + e log (q+1).

From the two given points on the cost function, the coefficients d and e were estimated by:

$$\hat{e} = \frac{c_1 q_2 - c_2 q_1}{q_2 \log(q_1+1) - q_1 \log(q_2+1)} \quad (16)$$

$$\hat{d} = \frac{1}{q_1}\left(c_1 - \hat{e}\log(q_1+1)\right) \quad (17)$$

Maximisation of the profit function

$$\pi = aq - bq^2 - dq - e\log(q+1) \quad (18)$$

yields the first order condition

[1] This convention and the corresponding conventions for Rules 5 and 6 may plausibly be considered to have introduced some bias against these rules in our final evaluation. For with *p* =21 sales are zero and profit negative.

$$\frac{d\pi}{dq} = a - 2bq - d - \frac{e}{q+1} = 0 \quad (19)$$

whence

$$q = \frac{(a-d-2b) \pm \sqrt{(a-d-2b)^2 + 8b(a-d-e)}}{4b} \quad (20)$$

The second order condition is

$$-2b + \frac{e}{(q+1)^2} < 0 \quad (21)$$

Various problem possibilities arise in this calculation, and they were handled by the conventions indicated below:

1) Double roots. If the single distinct q was negative, a price of 21 was charge. If the single distinct root was positive, but the second order condition was not satisfied, a price of 21 was charged, yielding zero sales.

2) If roots were complex, a price of 21 was charged.

3) If the roots were real and distinct, a price of 21 was charged either if both roots were negative or if the single positive root corresponded to a minimum.

If a positive q*, corresponding to a maximum of the profit function existed, the decision maker charged the price p = a-bq*.

Rule 6: Hyperbolic demand, quadratic cost. In Rule 6 the demand function is approximated by p = a-b/q, and its coefficients were estimated from the two observed points by:

$$\hat{b} = \frac{p_2 - p_1}{q_2 - q_1} \cdot q_1 q_2 \quad (22)$$

and

$$\hat{a} = p_1 - \frac{\hat{b}}{q_1} \quad (23)$$

The cost function was estimated as under Rule 4.

The profit function is now:

$$\pi = aq - b - dq - eq^2 \quad (24)$$

and, maximizing we obtain

$$\frac{d\pi}{dq} = a - d - 2eq = 0 \quad (25)$$

yielding

$$q = \frac{a-d}{2e} \quad (26)$$

and the second order condition is now e > 0.

If either the second-order condition was violated or the pseudo-optimal q was non-positive, the pseudo-optimal price p = 21 was chosen. Otherwise the decision-maker was taken to charge:

$$p = a - b\left(\frac{2e}{a-d}\right) \quad (27)$$

Appendix 2: Simulation runs results

Run 1	Max. Profit	Rule 1	Rule 2	Rule 3	Rule 4	Rule 5	Rule 6
Mean Profit realised: Expected Growth Option Wait Option KTL (Standard Deviation)	6827 (1210)	6403 (1273)	5233 (1535)	6110 (1473) 7739 (7046) 7236 (2213) 6047 (1473) 6002 (1500)	2854 (981) 7585 (1988) 3585 (2247) 6044 (1606) 5882 (1516)	2868 (935) 7674 (11603) 4574(2245) 5949 (1411) 5925 (1426)	277 (717) 1707 (1298) 96 (13615) 4006 (167) 3996 (308)
Mean profit/ Maximum Profit	1.000	0.938	0.832	0.945 1.133 1.065 0.885 0.879	0.418 1.111 0.525 0.885 0.861	0.420 1.124 0.669 0.873 0.867	0.040 0.250 0.014 0.586 0.585

Run 2	Max. Profit	Rule 1	Rule 2	Rule 3	Rule 4	Rule 5	Rule 6
Mean Profit realised: Expected Growth Option Wait Option KTL (Standard Deviation)	4913 (1406)	4113 (1333)	3053 (1852)	4145 (1539) 5788 (8724) 5144 (1953) 4414 (1172) 4302 (1174)	1909 (880) 3266 (7161) 4861 (1994) 4227 (1190) 4289 (1245)	1566 (1102) 27083(47333) 4438 (2021) 3946 (1625) 3914 (1657)	82 (522) 746 (11963) 309 (1886) 3954 (587) 3971 (309)
Mean profit/ Maximum Profit	1.000	0.837	0.624	0.843 1.178 1.047 0.898 0.875	0.388 0.664 0.989 0.860 0.872	0.318 5.512 0.903 0.803 0.796	0.016 0.151 0.062 0.804 0.808

Run 3	Max. Profit	Rule 1	Rule 2	Rule 3	Rule 4	Rule 5	Rule 6
Mean Profit realised: Expected Growth Option Wait Option KTL (Standard Deviation)	2779 (1417)	1047 (1670)	323 (2643)	1850 (1603) 4034 (13499) 2379 (1856) 3519 (1182) 3575 (1088)	806 (878) 2419 (11092) 2257 (1988) 3541 (1210) 3622 (1121)	557 (1116) 24745(54772) 2012 (1844) 2828 (1915) 2858 (1912)	-189 (632) 9381 (632) -277 (1149) 3876 (748) 3832 (2198)
Mean profit/ Maximum Profit	1.000	0.376	0.116	0.665 1.451 0.856 1.266 1.286	0.290 0.870 0.812 1.274 1.303	0.200 8.904 0.724 1.017 1.028	- 3.375 - 1.394 1.378

Appendix 2: Simulation runs results (continued)

Run 4	Max. Profit	Rule 1	Rule 2	Rule 3	Rule 4	Rule 5	Rule 6
Mean Profit realised: Expected Growth Option Wait Option KTL (Standard Deviation)	5866 (1244)	5125 (1322)	4041 (1835)	5042 (1410) 6649 (7203) 6064 (1957) 5016 (1416) 5019 (1428)	2449 (921) 6425 (7634) 3873 (2065) 5024 (1476) 5134 (1365)	2248 (1024) 10134(56008) 5885 (2181) 4756 (1491) 4736 (1553)	311 (874) 1023 (7850) 1298 (4368) 3964 (926) 4005 (1728)
Mean profit/ Maximum Profit	1.000	0.873	0.688	0.859 1.133 1.033 0.855 0.855	0.417 1.095 0.660 0.856 0.875	0.383 1.727 1.003 0.810 0.807	0.053 0.174 0.221 0.675 0.682

Run 5	Max. Profit	Rule 1	Rule 2	Rule 3	Rule 4	Rule 5	Rule 6
Mean Profit realised: Expected Growth Option Wait Option KTL (Standard Deviation)	4265 (1340)	2830 (1381)	1867 (2287)	3127 (1403) 5132 (8781) 4006 (1742) 3801 (1163) 3790 (1096)	1598 (2084) 2052 (16829) 4161 (2084) 3973 (1135) 3960 (1096)	1325 (977) 14350(12664) 3630 (2468) 3144 (1469) 3459 (1563)	3 (672) 1259 (12342) 162 (2424) 3915 (621) 3875 (677)
Mean profit/ Maximum Profit	1.000	0.663	0.437	0.733 1.203 0.939 0.891 0.888	0.374 0.481 0.975 0.931 0.927	0.310 3.364 0.851 0.737 0.811	0.000 0.295 0.037 0.917 0.908

Run 6	Max. Profit	Rule 1	Rule 2	Rule 3	Rule 4	Rule 5	Rule 6
Mean Profit realised: Expected Growth Option Wait Option KTL (Standard Deviation)	9246 (1245)	8790 (1294)	7516 (2016)	8828 (1201) 9442 (1300) 9720 (2271) 8793 (1274) 8877 (1251)	4474 (657) 4557 (1026) 9179 (2645) 8755 (1446) 8764 (1294)	4231 (894) 26168(36793) 9134 (2874) 8677 (1483) 8745 (1536)	440 (943) 126 (1342) 1161 (3873) 4059 (1478) 3994 (145)
Mean profit/ Maximum Profit	1.000	0.950	0.812	0.954 1.021 1.051 0.951 0.960	0.483 0.492 0.992 0.946 0.947	0.457 2.830 0.987 0.938 0.945	0.047 0.013 0.125 0.439 0.431