

Chapter 10.

Useful Heuristics

Benjamin Scheibehenne and Bettina von Helversen

Decision-making is one of the core tasks in project management. Traditionally, optimization methods have been developed to support managers in finding the best solutions. Alternatively, decisions can be based on simple rules of thumb, or heuristics. Even though simple heuristics only require little in the way of time and information, they have been shown to outperform optimization methods in complex decision tasks across a wide range of situations. This chapter outlines relevant decision heuristics commonly used, demonstrates situations in which they outperform more complex decision algorithms and explains why and when simple heuristics provide powerful decision tools.

Introduction

Imagine two managers who want to develop a real estate project. At some point during the planning phase it becomes crucial to predict the future sales prices of houses in a certain area. The first project leader, Mr O, approaches this task by searching exhaustively for all available pieces of information that he knows will influence the selling price, such as property tax, lot size, total living space, age of the house, number of bathrooms, and so on. Based on his past experience, he weighs all that information according to its importance and then integrates it to predict the selling price of each house, using some statistical software. The second manager, Mr F, makes a fast decision, relying on a simple strategy based on just one single piece of information that he regards as most important, such as total living space. Which of these two managers will make a more accurate forecast? Many people, researchers and laypersons alike, suppose that the outcome of a decision can be improved by (a) an exhaustive search for information, and the integration of many pieces of information, (b) having more time to think, and calculate possible outcomes, or (c) having more computational power, and the use of complex forecasting software or decision tools. This seems to indicate that Mr O will make the better decision. However, this chapter demonstrates that less can sometimes be more, and that a strategy relying on very few pieces of information, and quickly deriving a decision based on a simple algorithm, may well outperform more sophisticated, supposedly rational decision strategies.

To this end, the topic of “fast and frugal heuristics” is introduced (Gigerenzer et al, 1999). This investigates when and how simple decision strategies can be used to make sound decisions. The evidence shows that heuristics work well in many real-world situations that are characterized by uncertainty and scant information. Moreover, simple heuristics can be seen

to outperform more complex decision models. Before turning to the science of heuristics however, the status-quo of current decision research is examined, and shown to consist mainly of optimization models.

Optimization Models of Decision-Making

The unboundedly rational decision-maker

In his attempt to predict sale prices of houses, Mr O in the example above endorses a complex strategy that incorporates as much information as possible. In abstract terms, the underlying decision model could be described as an attempt to find the optimal solution, i.e. the solution that maximizes a given criterion measured by an objective function. In this case, the criterion was the probability of making a correct forecast; in other cases it might be the expected monetary gain, or the number of people who receive a benefit. Optimization models commonly try to find a solution by searching for all relevant information, weighing it according to its importance and combining it into an overall score. However, optimization implies that a well-defined criterion exists that can be used to calculate the function that is maximized. This further requires that the decision-maker has full knowledge of the decision task, can acquire all relevant information, and that the algorithm that leads to the optimal solution is known. It further requires that the decision-maker has the computational abilities and capacities to process all the information, or else has a decision tool readily available that fulfills these criteria respectively.

This “unboundedly rational” view on decision-making is highly unrealistic for almost every decision task in real-world environments, e.g. predicting future house prices. First of all, decision-makers are not omniscient: even for the simplest decision (let alone the more complex decision tasks faced by project managers) not all relevant information is known. On the contrary, to gather all relevant information can require painstaking amounts of effort, and may be expensive or impossible to achieve. Hence neither the time nor the money to search for and integrate the information may be affordable, or worth the possible increase in decision quality.

Furthermore, in many situations optimal solutions are not attainable. Even in well-defined decision problems with a limited number of options, such as chess, or the travelling salesman problem, optimal solutions are computationally intractable or NP-hard (Reddy, 1988). A problem is called NP-hard if its solution cannot be found in polynomial time. Roughly speaking, this means that no machine or mind can find the best solution in a sensible amount of time, such as a millennium.

In many cases, including the evaluation of major projects, there are multiple criteria to measure success, and the importance assigned to these criteria may vary dynamically over time, and between different stakeholders. Yet in cases where the criterion is ill-defined, when more than one goal exists (e.g. minimizing time, cost and risk of failure), or when stakeholders differ in their priorities, finding the optimal solution is utterly impossible (Gigerenzer, 2004).

Optimization under constraints

In an attempt to keep up with the gold-standard of optimality, some decision models try to take these restrictions into account. Aiming for “optimization under constraints”, these models

integrate search or information costs, by assuming that the decision-maker conducts a cost benefit analysis, and only continues to search for further information until the costs incurred in this search outweigh the benefits gained from the additional information. Even though these models try to impose constraints on decision-making, they often make even higher demands on the amount of information and computation necessary to conduct them. They not only assume that all information is available, but also that the cost of searching for the information, and its potential benefit, can be known or at least accurately estimated, to make an optimal decision. As such, they are simply another form of unbounded rationality (Arrow, 2004).

Decision tools widely used in project management, such as cost benefit analysis, or CBA (Pearce and Nash, 1981), multi-attributive utility theory (MAUT), multi-attributive value theory (MAVT), analytical hierarchy process (AHP), or decision trees, are examples of models that aim for optimization with or without constraints: they share a goal of finding the solution that maximizes (i.e. optimizes) the expected utility of stakeholders (Kiker et al, 2005). In line with the notion of optimization under constraints, these decision models further require a mathematical representation and integration of all possible information about risks and benefits. For instance, CBA tries to quantitatively evaluate the benefits, e.g. the expected profit or public utility, and the disadvantages, e.g. costs or risks involved in a decision, by expressing them in a single currency - money. The value of a decision is determined by the expected benefits minus the expected costs. Similarly, MAUT and AHP employ numerical scores to convey the merit of one option compared with others on a single “utility” scale.

Some assumptions made by these models can be problematic. For example, CBA makes the somewhat unrealistic assumption that objective prices and probabilities can be determined for any cost or benefit, including such things as human lives (Adams, 1995). MAUT and MAVT take into account that stakeholders differ in their preferences, and in the utilities that they assign to different outcomes. However, as with CBA, these models assume that such interests can be represented numerically, and traded-off against each other. Furthermore, all these methods are information-hungry and time-intensive.

The accuracy of these optimization methods hinges on the degree to which they manage to achieve an adequate representation of the decision problem. Yet, for several reasons, this representation can be difficult or impossible to attain. Project management decisions are often characterized by a high degree of foreseeable and unforeseeable uncertainty (De Meyer et al, 2002), which can make it too costly or impossible to gain a representation of the project environment that would allow the calculation of an optimal, or even close to optimal, course of action. Furthermore, in many projects, important decisions need to be made fast. For example, for many consumer goods the right timing is crucial, and in rapidly developing markets, technology product cycles may be very short. Here, an extensive and exhaustive search for information — which might be necessary to reach an accurate picture of the decision problem — can lead to a state of analysis-paralysis, where an opportunity is missed because of the sheer amount of information that needs to be processed.

Last but not least, even if an accurate representation could be reached quickly, more time and more information will not always lead to a better decision. Decisions in project management are commonly made in a dynamically changing environment. For instance, the importance assigned to benefits or the probability of risks is not stable, but will depend on the political and economic environment. This means that decision strategies need to be robust to perform well in changing environments, and able to adapt to new circumstances. Moreover, complex

decision algorithms are highly susceptible to change in their parameters, and can therefore lead to sub-optimal outcomes (Brighton and Todd, 2007; Todd and Gigerenzer, 2007).

To summarize, there is little doubt that optimization methods such as CBA or MAUT are useful tools for project managers, and can lead to informed and accurate decisions, particularly when the structure of a project is clear, and there is little uncertainty. However, in project decisions where uncertainty is high, risks and consequences difficult to estimate or even foresee, and the time available for analysis limited, they may feign a degree of certainty that is by no means substantiated by reality (De Meyer et al, 2002). In such situations, simple decision strategies can be a useful alternative, as they allow for fast and robust decisions in dynamically changing environments.

Heuristics for Decision-Making

Most decision scientists acknowledge that heuristics are widely used, and yet there is controversy over why people employ heuristics. Some scholars see heuristics as a crutch that is necessary because the memory and processing capacities of the human mind are severely limited. This impedes the use of more complex strategies, and leads to the failure to consider all important information. From this perspective, heuristics are often seen as a flaw that will lead to second-best outcomes, systematic biases, and a distorted evaluation of confidence in the decision (Kiker et al, 2005; McDaniels et al, 1999; Tversky and Kahneman, 1974 and 1979). As a consequence, decision-makers would be well-advised not to use them, but to rely on more complex strategies, or decision support tools instead.

So are heuristics only for those who cannot appreciate the blessings of modern technology, and thus remain irrational and biased? Or can there also be such a thing as a *useful* heuristic? Intuitively, one might think that a decision based on a simple strategy, that uses little information, cannot be as sound as a decision in which many pieces of information have been evaluated, weighted, and integrated. Yet, as it turns out, the chances are that this intuition is wrong.

The research on “fast and frugal heuristics” (Gigerenzer et al, 1999) embraces heuristics as adaptive responses to our environment, and their ability to ignore information as a strong virtue. Rather than considering the use of heuristics a flaw of human decision-making, Gigerenzer argues that heuristics are decision mechanisms which evolved to enable humans to make quick and accurate decisions in uncertain environments. As the proverb ‘time is money’ indicates, our world is competitive. Any time spent on decision-making keeps us away from other activities which could increase our chances to outperform our competitors. Because it takes time to search for, process, and assess information, decision mechanisms which rely on scant information or computation have the advantage of speed. Further, simple mechanisms can be more accurate than complex mechanisms, because they are more robust. Robustness is an important feature of decision mechanisms in uncertain and dynamic environments. Because the same decision is rarely encountered twice, and decision environments are unstable, decision mechanisms need to adapt well to new circumstances (Brighton and Todd, 2007; Todd and Gigerenzer, 2007). Lastly, simple heuristics are highly transparent decision strategies. Because of their simplicity, they can be easily communicated, and are thus more likely to gain public acceptance.

According to Gigerenzer et al (1999), the mind consists of an “adaptive toolbox” of fast and frugal heuristics (Gigerenzer and Selten, 2001), which evolved to solve certain decision problems, such as a choice between two or more options. These heuristics are adapted to specific information structures, which they exploit to allow for fast and accurate decisions.

The research on fast and frugal heuristics pursues a descriptive and a normative goal. It is descriptive, in that it tries to specify which heuristics decision-makers employ in a given situation, and it is normative, by showing that these heuristics are, in fact, adaptive solutions to specific decision environments.

Some of these fast and frugal heuristics are described in the following sections. It will be seen that they are widely adopted, and can be as good as – and sometimes even better than – “optimal” decision strategies.

Relying on Just One Good Reason: The Take-the-Best Heuristic

In the introductory example, Mr F only relied on one good reason to predict the selling price of houses. He did not weigh or trade-off different pieces of information, but chose the best option for the aspect that he regarded as most important. This particular simple heuristic is called take-the-best because it gambles on the most valid piece of information (Gigerenzer and Goldstein, 1996). It belongs to a class of heuristics known as “lexicographic” (LEX), because they resemble a lexicon, in which words are strictly sorted by the order of their first letters (Payne et al, 1993).

A typical LEX strategy works as follows: In the first stage of the information search, the single most important aspect (i.e. the most valid cue) is taken into consideration. If one option is shown to be better than all the others with regard to this cue, the search is stopped and the option with the highest cue value is chosen. If several options are equally good for the most valid cue, the next best aspect is considered as a tie-breaker until a decision can be made (Figure 1). In many practical contexts, options are never completely similar with regard to any given aspect. In these cases, a threshold can be set that determines the point at which two options are still regarded as being equal.

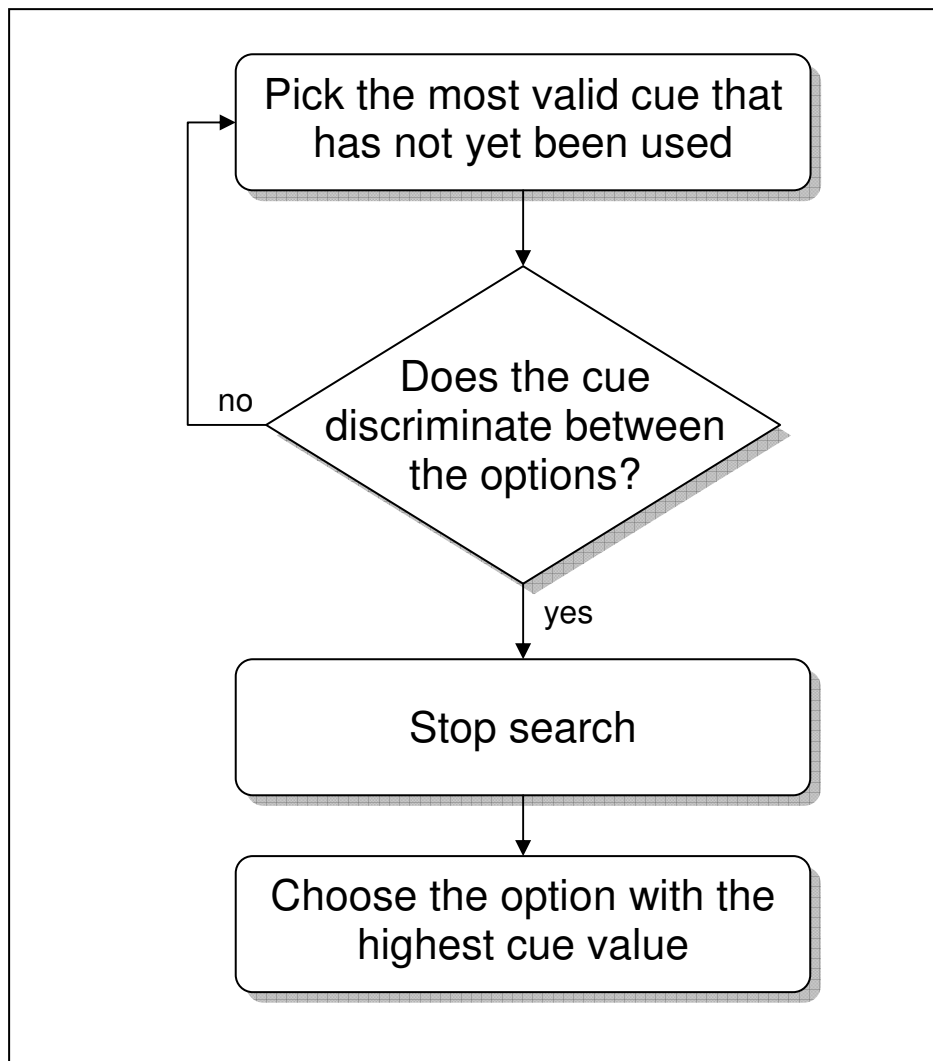


Figure 1 Flow chart of a simple lexicographic decision rule

LEX is a non-compensatory heuristic, which means that an advantage on an important aspect cannot be compensated by one (or a combination of) less important aspects. Because less important aspects are seldom needed to derive a decision, LEX requires little information search. But how good is LEX compared with other, more complex decision strategies?

Keeney and Raiffa (1993) argued that one-reason decision-making “is more widely adopted in practice than it deserves to be,” is “naively simple,” and “will rarely pass a test of reasonableness” (p.77-78). However, while they never provided such a test, others did.

In a series of computer simulations based on 20 real-world data sets, including the prediction of house-sale prices, Czerlinski et al (1999) compared the LEX heuristic with information-hungry forecasting methods, such as multiple linear regression. In the case of fitting the data sets, the more flexible multiple regression analysis led to the highest percentage of correct decisions. However, if the criterion for comparison was a more realistic prediction task, where the models had to predict the future, rather than merely post-predicting the past, they found to their surprise that LEX repeatedly outperformed its competitors.

The effectiveness of simple decision heuristics is not limited to the data sets analyzed by Czerlinski et al, or comparisons to multiple regression analysis. Brighton (2006) showed that, across a number of inductive inference tasks, LEX also outperformed a wide selection of

state-of-the-art machine learning algorithms, including neural networks, exemplar models, and elaborate classification trees (e.g. CART). As a performance measure, Brighton used the models' ability to accurately predict new data in real-world environments, as well as the "minimal description length" criterion (Pitt et al, 2002), which indicates how well a model can compress the data. With regard to both criteria, in most situations LEX did as well as, or even better than its allegedly more powerful competitors.

Relying on the First Reason that Comes to Mind: The Take-the-First Heuristic

The usefulness of simple heuristics is not constrained to computer simulations and macro-economic data, but can be applied to many real-life contexts, including the domain of professional sports. Imagine a handball player in a fierce game who has to decide on the next move. As handball is a fast sport, decisions have to be made within fractions of a second. Wrong decisions may cause the loss of the ball, and sometimes even the whole game. In such a situation, the classical rational choice model would predict that having more time to think before passing the ball (or being able to process more information in the same time) would give players a competitive advantage, because they have more time to evaluate their options, estimate the probabilities that a certain move will be successful and predict the behaviour of the opponent. Yet, it turns out that this is not the case. In a series of studies on handball players, Johnson and Raab (2003) found that those players who had little time to think made better decisions than those who were given more time. Similarly, Beilock et al (2004) showed that professional golfers were more accurate if they had less time to consider how to approach the task.

After thorough analyses of the data, it appeared that players relied on the "take-the-first" heuristic, i.e. when time-pressure was tight, they chose the first move or pass that came to mind. As it turned out, for these highly experienced players, the options that came up first were usually the better ones. Even though more time enabled them to generate more options, these were shown to decrease in quality. Eventually, the players with more time to consider ended up with more low-quality options on the table, which, in turn, increased the chances that one of the low-quality options was selected.

Equal Spread among All Options: The 1/N Heuristic

It could be argued that the success of simple heuristics might only apply to small or somewhat less important decisions, in which people avoid the additional effort required by a more thorough strategy. This argument can be countered by studying the more "important" case of financial asset allocation. To find out how people in the real world go about investing their money, Huberman and Jiang (2006), analyzed the records of more than half a million people who participated in pension plans (401(k) plans). They found that, when deciding in which funds to invest, the majority of people chose a small number of funds, and then allocated their contributions evenly across them (Benartzi and Thaler, 2001; Loomes et al, 1991). This strategy is known as the 1/N heuristic, which can, indeed, be traced back to the 4th century, where Rabbi Isaac Bar Aha in the Talmud recommended splitting one's wealth equally among several investments.

From the perspective of neoclassical economic theory, the 1/N heuristic constitutes poor reasoning, sometimes mocked as "couch potato investments" or "coward portfolios", because investors ought rather to base their decisions on sophisticated statistical methods, such as probabilistic scenario analyses, that aim to optimize the mean variance, or the portfolio risk-return profile (Huberman and Jiang, 2006).

However, in a thorough comparison of fourteen optimization models for portfolio choice across several empirical data sets, De Miguel et al (2007) found that none was consistently better than the apparently naïve 1/N heuristic, in terms of Sharpe ratio, certainty-equivalent return, or turnover. They showed that in order for the “optimal” portfolio strategies to achieve a higher certainty-equivalent return than the 1/N heuristic, the optimal models would need a portfolio with only 25 assets to have 291 years’ worth of stock-market data for parameter estimation. For a portfolio with 50 assets, this period increases to more than 500 years, which is in sharp contrast to the common practice in which the parameters of these models are typically estimated using only 5 to 10 years of data (De Miguel et al 2007). Thus, while the optimal models might theoretically lead to a superior outcome, for any real-world situation, in which people commonly do not have the patience to wait for 500 years, applying the 1/N heuristic is a sensible alternative.

Given these results, it is little surprise that the 1/N heuristic is not only used by laypeople, but also by designated experts in the field. The Nobel laureate, Harry Markowitz, who developed a model for optimal portfolio selection (Markowitz, 1952), reported that he used the 1/N heuristic for his own, private investments (Gigerenzer, 2008). Moreover, the tendency to equally spread one’s investments is not confined to financial decisions, and reliance on the 1/N heuristic can be observed in many other circumstances. For example, parents usually try to invest equally in their children (Hertwig et al, 2002), and across many cultures, equal splitting is a widely accepted norm for resource allocation (Leman et al, forthcoming).

Simple Heuristics in Project Management

In a thorough analysis of the role of uncertainty in project management, Pich et al (2002) acknowledge that project teams often use heuristics to generate their policies, suggesting that the proliferation of heuristics is also apparent in management and business contexts, especially when it comes to project scheduling problems (Davis and Patterson, 1975; Russell, 1986; Hartman and Kolisch, 2000).

One recent and well-documented example of how heuristics can be successfully applied in project management stems from a study published in the renowned journal *Management Science* by Åstebro and Elhedhli (2006). They provided evidence that experts’ decisions were essentially made by the use of a heuristic strategy, and that the use of this simple heuristic outperformed a sophisticated log-linear multiple regression analysis in predicting the future commercial success of R&D projects at an early stage.

In their study, Åstebro and Elhedhli used records from 561 randomly selected R&D projects. Outcomes of R&D projects are notoriously uncertain, while at the same time, the potential feedback is patchy, and decision-relevant information is not easily quantified – hence, they chose a situation where accurate predictions are especially hard to come by. For each project, they acquired information indicating whether the product successfully reached the market (the criterion that had to be predicted), as well as 37 independent pieces of information (e.g. cost of production, tooling costs, existing competition, etc.), that could already be assessed at an early stage of the projects, and thus could be utilized as cues for forecasting. Based on qualitative interviews, Åstebro and Elhedhli found that, in order to predict project success, many expert analysts used a strategy described as a “non-compensatory tallying heuristic”, which consisted of a search and a decision phase. Firstly, the experts evaluated the projects on 37 different criteria or cues, including aspects such as technical feasibility, expected costs, and market development. For each cue, the experts rated if the project would do well; if it would do poorly; or if it would be critically flawed. Next, they added up the number of “good” and “bad” cues separately. If no critical flaws existed, if the number of good cues exceeded a

fixed threshold g , and the number of bad cues was lower than a fixed threshold b , they predicted a success; if not, they predicted a failure. It turned out that this heuristic outperformed a competing logistic regression model which incorporated all available cues in out-of-sample prediction. It also did better than a stepwise log-linear model using backwards variable elimination.

The tallying model that mimicked experts' forecasting rules still used 33 out of the 37 possible cues. Interestingly, further analyses by Åstebro and Elhedhli showed that experts actually used considerably more information than necessary, and that they could have improved their predictive accuracy by using only 21 cues. Zacharakis and Meyer (2000) showed a similar result, where venture capitalists improved their predictions when provided with less information about the venture.

Similar results, showing that less information and less computational complexity can lead to better outcomes were reported as long ago as 1979, by Makridakis and Hibon, who compared the accuracy of several statistical forecasting models in predicting future data ("out-of-sample" accuracy). The comparison was based on 111 real-life time series, covering a wide range of contexts, including business, industry, and macro-economic data. What they found was that simple forecasting methods, e.g. calculating an ordinary moving average, repeatedly outperformed sophisticated ones, e.g. ARIMA models, or multiple linear regression. Even though they were not the first to find such results (Reid, 1969 and 1975; Newbold and Granger, 1974; Dawes, 1979), their findings provoked strong objections from statisticians and other decision scientists. To meet these objections, Makridakis et al (1982) launched a number of "M-Competitions", in which they invited their critics to let their models compete in a new set of time series. While the competing models were even more sophisticated than before (e.g. decomposition models, expert systems and neural networks), the results were strikingly similar to those in the earlier study. Complex methods did not provide more accurate forecasts than simpler ones (Makridakis and Hibon, 2000).

This is a small sample from countless well-documented cases in which laypeople, as well as experts, employ simple heuristics to accomplish their goals across a wide range of situations (Gigerenzer et al, 1999; Gigerenzer, 2004b). This suggests that (a) heuristics are not an exception, but rather the rule when it comes to decision-making in dynamic real-world environments, that are characterized by uncertainty and information complexity; and that (b) experts who rely on simple heuristics can be as good as or even better at predicting uncertain outcomes than machines relying on complex statistical models.

Why do project managers need to know how people reason?

With regard to the descriptive aspect of heuristics research, it might be argued that knowledge about human decision-making strategies in real-world situations is a purely academic question that lacks practical importance. On the contrary, it can be argued that the success of many projects crucially depends on project managers' understanding of how people actually reason, and that, in many situations, people employ rather fast and frugal heuristics to derive their decisions.

In 1998, the German government released a law that allowed private households to freely choose their electricity provider. Similar to other attempts of market liberalization, one of the underlying rationales of this policy was to increase competition and lower prices. Yet a couple of years later it turned out that only very few consumers had switched their provider, even though most of them could have saved money by switching. From the perspective of humans as fully rational agents, who strive towards utility-maximizing, a convincing explanation for

this apparent inertia to change was that people simply did not know that they had a choice, or that they did not know the alternatives. Consequently, in 2002, a major German electricity producer launched a large-scale advertising campaign worth €22M., informing private households about their opportunity to switch. To their surprise, by the end of the campaign as few as 1,100 customers had made use of this offer, resulting in a net cost of €20,000 for each new customer – rather an expensive advertising campaign.

From the perspective of decision research, one explanation for this miscarriage is that managers had the wrong representation of humans as utility-maximizers who trade-off costs and benefits, and who would switch their provider as soon as there is a marginal monetary incentive to do so.

The assumption that people trade-off aspects against each other, and that they choose options with the highest net outcome, is not unique to German managers. When it comes to tools used to analyze and describe human decision-making, it is the rule rather than the exception. One prominent representative of such a tool, which remains to be widely used, is conjoint analysis. The term “conjoint analysis” stems from “consider-jointly”, and describes a collection of sophisticated methods that assume decision-makers weigh the pros and cons (i.e. the risks and benefits) of each expected outcome, in order to derive an overall expected utility for each option. This feature of conjoint analysis is similar to the notion of multi-attributive utility theory (MAUT), and is shared by the method of cost benefit analysis (CBA), where money, rather than expected utility, depicts the common currency.

In sharp contrast to this notion, a growing body of literature shows that human reasoning might better be described and predicted by heuristic processes. In a review of 45 studies that investigated people’s decision strategies in a variety of contexts, including choices between apartments, microwaves, and birth control methods, Ford et al (1989) found that decision-makers often evaluate surprisingly small amounts of information, and in many situations do not trade-off the pros and cons of the options, but rather rely on a lexicographic heuristic.

Along the same lines, a recent study by Yee et al (2007) found that the simple LEX strategy predicted people’s decisions better than a state-of-the-art conjoint analysis using highly parameterized linear models and Bayesian estimation algorithms. In their experiment, participants chose from an assortment of smart phones that differed in design, functionality, operating system, and other attributes. The results were subsequently confirmed by another group of researchers, who studied choices between laptop computers (Kohli and Jedidi, 2007).

Insofar as humans tend to make decisions without considering all information, and without making many trade-offs, decision tools which assume a fully rational and compensatory decision process, can grossly misrepresent the actual decisions and preferences of the target group. This can lead to wrong predictions of market potential, as is illustrated by the example above.

Moreover, decision-makers can also be influenced by procedures to elicit their preferences. For instance, Wilson and Schooler (1991) and Wilson et al (1993) showed that people chose different options depending on whether they were instructed to make a decision based on their intuition, or on more informed consideration. More importantly, after two weeks, those who had relied on their gut feelings were more satisfied with their choices than those who had made a considered choice. This indicates that procedures such as MAUT, that enforce a compensatory decision strategy, might lead to inferior and less satisfying decision outcomes.

Why do heuristics work?

It has been shown that decision-makers use heuristics, and that these can outperform even vastly complex and information-hungry decision tools. Knowing this, the next challenge is to explain why and in which situations heuristics are useful. The key to answer these questions are two principles that are outlined in more detail below: domain-specificity and robustness.

Heuristics are domain specific. Heuristics are capable of ignoring a lot of “noise”, and thus make robust predictions, because they can exploit certain patterns of information in the environment. For example, when trying to predict the outcome of the Men’s Singles competition at Wimbledon, one could employ the “recognition-heuristic” (Goldstein and Gigerenzer, 2002). For predicting Wimbledon outcomes, this heuristic works as follows: if a player whose name is recognized competes against an unknown player, the prediction is that the known player will win the match. This, essentially simple, heuristic yields a surprisingly high accuracy, that can be as good as expert predictions, or those based on official ATP rankings (Scheibehenne and Bröder, 2007; Serwe and Frings, 2006). As the heuristic relies on partial ignorance, it can only be used for those cases in which one of the players is unknown. If neither player is known, it does not apply. The success of the recognition heuristic depends on certain information structure in the environment, namely that there is a systematic relationship between name recognition and success. Here, successful players gain more media attention in the time before the competition, and those players also have a higher probability of winning a match. The degree to which name recognition is linked to a certain criterion (e.g. success at Wimbledon) is known as “recognition validity”. The recognition heuristic is domain specific, in that it works well in environments where the recognition validity is high.

There are a number of studies showing that recognition validity is surprisingly high across a wide range of situations, yet it is important to stress that it is not always the case (Pohl, 2006). This leads to an important aspect about why and when heuristics work. Their success depends on how well they fit the structure of the environment. Thus, heuristics are not rational per se; rather they are “ecologically rational”. The Nobel laureate, Herbert Simon (1990), illustrated this principle using the analogy of a pair of scissors with two blades. One blade depicts the decision strategy (i.e. the heuristic) while the other depicts the structure of the environment (i.e. the situation or circumstances). The scissors only cut when the two blades match up.

To illustrate this principle, the findings by Johnson and Raab (2003) showed that time-pressure led to better decisions for highly trained handball players. These players could benefit from the take-the-first heuristic, because the options that first came to their minds were usually the better ones, and the quality of additional options decreased over time. With this in mind, and in line with the idea of ecological rationality, Johnson and Raab further showed that, unlike experienced handball players, unskilled players did not benefit from time-pressure. For them, probably due to their lack of experience, good options did not necessarily leap out, but rather had to be searched for. In this case, having more time increased their chances of finding a good option. These results suggest that a heuristic such as take-the-first is not universally recommended, but rather depends on a certain environment - in this case, the individual player experience.

Similarly, with respect to the 1/N heuristic outlined earlier, DeMiguel et al (2006) identified circumstances in which the heuristic underperformed against optimization models, namely in situations where the number of investible assets is small, the assets differ widely in their return, and the estimation window is long.

These examples show that the usefulness of a heuristic is context-dependent, and that there is no such thing as a heuristic that is useful for everyone, or in every situation. Rather, heuristics are tools that are tailored specifically to a given situation or problem. Just as a screwdriver or pair of pliers is made to solve a specific task, and fits into its environment, i.e. the head of the screw, each heuristic evolves to solve a specific decision task, and achieves this by exploiting the structure of information in the environment. To continue with the metaphor of heuristics as specific tools, the repertoire of heuristics that decision-makers can choose from is referred to as the “adaptive toolbox” (Gigerenzer and Selten, 2001).

It might be asked, though, how people go about picking the right tool from the toolbox? There are many possible ways. Probably the most important one is that of learning through feedback. For example, in an experiment by Rieskamp and Otto (2006), minimal outcome feedback was apparently sufficient to learn which heuristic worked best in a given situation. Alternatively, to select the right heuristics from the toolbox, decision-makers can also resort to imitation of the successful (Garcia-Retamero et al, 2006), explicit instructions, or advice from more experienced people, to name but a few.

Heuristics are robust. In hindsight, everything else being equal, a more complex model will always outperform a simple one. Yet for most real-world situations, what matters is out-of-sample accuracy - the prediction of the future. For example, after a project is finished, it may be easy to identify the reasons and the chain of events that eventually led to its success or failure. While this might be interesting information in itself, what one really wants to know is if these reasons will still apply for predicting the success of future projects. However, predicting the future is much harder than merely giving ex-post explanations of things that have already happened.

When attempting to predict upcoming future developments, available information can be divided into two parts: one part provides relevant data, the other part is irrelevant, or “random noise”, that should best be ignored. The tricky part is to decide which particular piece of information is relevant, and which is not. In general, the more uncertain the environment, the less information can be used to predict the future, and the more data should be ignored (Gigerenzer, 2008). For example, when predicting the outcome of the stock market in the future, most past events are irrelevant, because they are already fed into the current stock prices. In this situation, a sophisticated model that incorporates lots of this past information will be quite capable of giving an accurate description of available data, but it will perform poorly in predicting an uncertain future, where new events will occur that do not necessarily resemble the past. This discrepancy between the ability to describe readily available data, and to predict new data is known as “overfitting”.

The reason why simple models often work better than more complex ones is because they ignore less important pieces of information, and tend to focus on just a few, valid cues. For example, heuristics like LEX and take-the-first only use one very important cue, and ignore all the remaining information. Simple heuristics can reduce the risk of overfitting, a property sometimes referred to as “robustness” and one that is especially valuable in uncertain environments.

Intuitively, one might think that the more uncertain an environment, the more information one should gather, and yet, as outlined above, a large body of evidence in such diverse fields as artificial intelligence, linguistics, psychology, economics, and operations research suggests that it is precisely the other way round: the more uncertain the outcome, the less information should be employed to make a choice.

Conclusion

In uncertain and dynamic environments, and for many relevant project decisions, an optimization approach is often difficult, and sometimes nigh impossible. In such situations, heuristics provide a feasible way to make decisions. Contrary to the common view of heuristics as second-best solutions that people only use because of their limited information-processing capabilities, the research on fast and frugal heuristics has provided substantial evidence that heuristics often achieve an astonishingly high performance, using just a fraction of the time and the amount of information required by standard decision strategies. Heuristics achieve this success for two reasons: firstly, they are ecologically rational, i.e. they are adapted to a specific task environment; secondly, they are robust, because they ignore much information, which makes them less vulnerable to random noise. By focusing only on relevant pieces of information, heuristics are well able to adapt to new situations.

To summarize, even though research on simple heuristics continues to emerge, the insights gained thus far already disclose a promising approach towards understanding and improving decision-making in project management, for instance, when allocating resources, or in predicting future outcomes at an early stage.

References

Adams J (1995). *Risk*. UCL Press: London.

Arrow K (2004). Is bounded rationality unbounded rational? Some ruminations. In: Augier M and March J (eds). *Models Of A Man: Essays in Memory of Herbert A. Simon*. MIT Press: Cambridge, MA.

Astebro T and Elhedhli S (2006). The effectiveness of simple decision heuristics: Forecasting commercial success for early-stage ventures. *Management Science* 52: 395-409.

Beilock S, Berenthal B, Mccoy A and Carr T (2004). Haste does not always make waste: expertise, direction of attention, and speed versus accuracy in performing sensorimotor skills. *Psychonomic Bulletin & Review* 11: 373-379.

Benartzi S and Thaler R (2001). Naive diversification strategies in defined contribution saving plans. *American Economic Review* 91: 79-98.

Brighton H (2006). Robust inference with simple cognitive models. In: Lebiere C and Wray B (eds). *Between a rock and a hard place: Cognitive science principles meet AI-hard problems. Papers from the AAAI Spring Symposium*. AAAI Press: Menlo Park, CA.

Brighton H and Todd P (2007). Situating rationality: Ecologically rational decision making with simple heuristics. In: Robbins P and Ayedede M (eds). *Cambridge Handbook Of Situated Cognition*. Cambridge University Press: Cambridge.

Czerlinski J, Gigerenzer G and Goldstein D (1999a). How good are simple heuristics? In: Gigerenzer P and The ABC Research Group (ed). *Simple heuristics that make us smart*. Oxford University Press: New York.

Davis E and Patterson J (1975). A comparison of heuristic and optimum solutions in resource-constrained project scheduling. *Management Science* 21: 944-955.

Dawes R (1979). The robust beauty of improper linear models in decision making. *American Psychologist* 34: 571-582.

De Meyer A, Loch C and Pich M (2002). Managing Project Uncertainty: From Variation to Chaos. *Mit Sloan Management Review* 43: 60-67.

DeMiguel V, Garlappi L, and Uppal R (2007). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *Review of Financial Studies*, doi:10.1093/rfs/hhm075

Ford J et al (1989). Process tracing methods: Contributions, problems, and neglected research questions. *Organizational Behavior and Human Decision Processes* 43: 75-117.

Garcia-Retamero R, Takezawa M and Gigerenzer G (2006). How to learn good cue orders: When social learning benefits simple heuristics. In: Sun R and Miyake N (eds). *Proceedings of the 28th annual conference of the Cognitive Science Society*. Erlbaum Mahwah: New Jersey.

- Gigerenzer G (2004). Striking a blow for sanity in theories of rationality. In: Augier M and March J (eds). *Models of a man: Essays in memory of Herbert A. Simon*. MIT Press: Cambridge, MA.
- Gigerenzer G (2004b). Fast and frugal heuristics: The tools of bounded rationality. In: Koehler D and Harvey N (eds). *Blackwell handbook of judgment and decision making*. Blackwell Publishing: Oxford.
- Gigerenzer G (2008). Why Heuristics work. *Perspectives on Psychological Science*.
- Gigerenzer G and Goldstein D (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review* 103: 650-669.
- Gigerenzer G and Selten R (2001). *Bounded rationality: The adaptive toolbox*. MIT Press: Cambridge, MA.
- Gigerenzer G, Todd P and The ABC Research Group (1999). *Simple Heuristics that make us smart*. Oxford University Press: New York.
- Goldstein D and Gigerenzer G (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review* 109: 75-90.
- Hartmann S and Kolisch R (2000). Experimental evaluation of state-of-the-art heuristics for the resource-constrained project scheduling problem. *European Journal of Operational Research* 127: 394-407.
- Hertwig R, Davis J and Sulloway F (2002). Parental investment: How an equity motive can produce inequality. *Psychological Bulletin* 128: 728-745.
- Huberman G and Jiang W (2006). Offering versus choice in 401(k) plans: Equity exposure and number of funds. *Journal of Finance* 61: 763-801.
- Johnson J and Raab M (2003). Take the first: Option-generation and resulting choices. *Organizational Behavior and Human Decision Processes* 91: 215-229.
- Keeney R and Raiffa H (1993). *Decisions with multiple objectives*. Cambridge University Press: Cambridge.
- Kiker G et al (2005). Application of multicriteria decision analysis in environmental decision making. *Integrated Environmental Assessment and Management* 1: 95-108.
- Kohli R and Jedidi K (2007). Representation and inference of lexicographic preference models and their variants. *Marketing Science* 26: 380-399.
- Kurz-Milcke E and Gigerenzer G (2007). Heuristic decision making. *Marketing: Journal of Research and Management* 3: 48-56.
- Leman P, Keller M, Takezawa M and Gummerum M (in press). Children's and adolescents' decisions about sharing money with others. *Social Development*.
- Loomes G, Starmer C and Sugden R (1991). Observing violations of transitivity by experimental methods. *Econometrica* 59: 425-439.

- Makridakis S et al (1982). The accuracy of extrapolation (time series) methods: Results of a forecasting competition. *Journal of Forecasting* 1: 111-153.
- Makridakis S and Hibon M (1979). Accuracy of forecasting: An empirical investigation. *Journal of the Royal Statistical Society A* 142: 97-145.
- Makridakis, S. & Hibon, M. (2000) The M3-competition: Results, conclusions and implications. *International Journal of Forecasting*, 16, 451-476.
- Markowitz H (1952). Portfolio Selection. *Journal of Finance* 7: 77-91.
- McDaniels T, Gregory R and Fields D. (1999). Democratizing risk management: Successful public involvement in local water management decisions. *Risk Analysis* 19: 497-510.
- Newbold P and Granger C (1974). Experience with forecasting univariate time series and the combination of forecasts. *Journal of the Royal Statistical Society* 137: 131-165.
- Payne J, Bettman J and Johnson E (1993). *The adaptive decision maker*. Cambridge University Press: New York, NY, USA.
- Pearce D and Nash C (1981). *The Social Appraisal of Projects*. Macmillan: London.
- Pich M, Loch C and De Meyer A (2002). On uncertainty, ambiguity, and complexity in project management. *Management Science* 48: 1008-1023.
- Pitt M, Myung I and Zhang S (2002). Toward a method of selecting among computational models of cognition. *Psychological Review* 109: 472-491.
- Pohl R (2006). Empirical tests of the recognition heuristic. *Journal of Behavioral Decision Making* 19: 251-271.
- Reddy R (1988). Foundations and grand challenges of artificial intelligence: AAAI Presidential Address. *AI Magazine*.
- Reid D (1969). A comparative study of time series prediction techniques on economic data. PhD Thesis. *Department of Mathematics*. University of Nottingham: Nottingham.
- Reid D (1975). A review of short term projection techniques. In: Gordon H (ed). *Practical aspects of forecasting*. Operational Research Society: London.
- Rieskamp J and Otto P (2006). SSL: A theory of how people learn to select strategies. *Journal of Experimental Psychology-General* 135: 207-236.
- Russell R (1986). A comparison of heuristics for scheduling projects with cash flows and resource restrictions. *Management Science* 32: 1291-1300.
- Scheibehenne B and Bröder A (2007). Predicting Wimbledon 2005 tennis results by mere player name recognition. *International Journal of Forecasting* 23: 415-426.
- Serwe S and Frings C (2006). Who will win Wimbledon? The recognition heuristic in predicting sports events. *Journal of Behavioral Decision Making* 19: 321-332.
- Simon H (1990). Invariants of human behavior. *Annual Review of Psychology* 41: 1-19.

Todd P and Gigerenzer G (2007). Environments that make us smart: Ecological rationality. *Current Directions in Psychological Science* 16: 167–171.

Tversky A and Kahneman D (1974). Judgment under uncertainty: Heuristics and biases. *Science* 185:1124-1131.

Tversky A and Kahneman D (1979). Prospect theory: An analysis of decision under risk. *Econometrica* 47: 263-91.

Wilson T et al (1993). Introspecting about reasons can reduce post-choice satisfaction. *Personality and Social Psychology Bulletin* 19: 331-339.

Wilson T and Schooler J (1991). Thinking too much: Introspection can reduce the quality of preferences and decisions. *Journal of Personality and Social Psychology* 60: 181-192.

Yee M, Dahan E, Hauser J and Orlin J (2007). Greedoid-based noncompensatory inference. *Management Science* 26: 532-549.

Zacharakis A and Meyer G (2000). The potential of actuarial decision models: Can they improve the venture capital investment decision? *Journal of Business Venturing* 15: 323-346.

1/N Heuristic;7
 adaptive toolbox;4
 analytical hierarchy process;3
 Bayesian estimation algorithms;10
 complex decision algorithms;4
 conjoint analysis;10
 cost benefit analysis;3
 coward portfolios;7
 decision trees;3
 decomposition models;9
 domain specific;11
 expected utility of stakeholders;3
 expert systems;9
 fast and frugal heuristics;1; 4
 gut feelings;10
 heuristic strategy;8
 heuristics;4
 heuristics research;9
 human decision-making strategies;9
 human reasoning;10
 information structures;4
 LEX strategy;5
 lexicographic;5
 M-Competitions;9
 multi-attributive utility theory;3
 multi-attributive value theory;3
 neural networks;9
 non-compensatory heuristic;6
 non-compensatory tallying heuristic;8
 Optimization under constraints;2
 overfitting;12
 parameterized linear models;10
 Robust heuristics;12
 Robustness;4
 simple heuristics;1
 take-the-best;5
 Take-the-First Heuristic;7
 unbounded rationality;3
 unboundedly rational decision-maker;2
 useful heuristic;4