Imagination Lab

Organizational Heuristics: More than Simple Rules

Working Paper 13

March 2002

DR DAVID OLIVER*
Research Fellow

DR JOHAN ROOS*
Director

*Imagination Lab Foundation - Rue Marterey 5 - 1005 Lausanne - Switzerland
Tel +41 21 321 55 44 - Fax +41 21 321 55 45 - www.imagilab.org

Abstract

This paper explores the nature of heuristics in organizations operating in complex environments. The authors review the current academic literature on heuristics, and then conduct interpretive case studies to explore how two organizations used heuristics in practice. These findings lead to the development of four dimensions along which heuristics vary: scope, focus, development and depth. These dimensions form the basis for four formal propositions concerning how organizational heuristics may be most effectively developed and used. It is proposed that effective heuristics should be general, internally and externally balanced, developed in a cyclical manner, and grounded in emotion as well as pure rationality.

Organizational Heuristics: More than Simple Rules

Introduction

It has recently been proposed that developing strategies based on "simple rules" is an effective way of coping with complex business environments (Eisenhardt 2001). Such an approach is implicitly heuristic[1], that is, it is based on using "rules of thumb" to find acceptable solutions to complex problems. Heuristics have long been the focus of a significant literature in social psychology (Kahneman and Tversky 1973) and have more recently been entering the organizational science realm (e.g. Gigenrenzer and Todd 2000). In the face of all the emergent and unplanned events experienced by firms today, how can organizational members develop effective rules of thumb?

It is the purpose of this paper to explore the nature of heuristics in organizations operating in complex environments. We do this by studying how two organizations have used heuristics, developing four dimensions that categorize these heuristics, and proposing four formal propositions concerning how they might be developed and used by others.

What are Heuristics?

At least since the Enlightenment, the idealized form of decision-making has rested on key assumptions of *rationality*, which is defined by Allison (1971: 30) as: "consistent, value-maximizing choice within specified constraints." Rationality has been assumed to be important because "if a person acts rationally, his behaviour can be *fully explained* in terms of the goals he is trying to achieve" (Harsanyi 1966, cited in Allison 1971:31). The classical "economic man" or "rational actor" faces a well-defined problem and has adequate skill, time and resources to select the well-defined alternative (about which he or she has full

[1] The origin of the word heuristics can be traced to the Greek word "eurisco", meaning "I discover."

information) that will lead to the highest payoff function. However, in the 1950s, Herbert Simon pointed out that these conditions rarely—if ever—exist. Simon proposed that people do not seek out a perfect solution to a given problem in a rational manner. Instead, they seek a solution that they consider ”good enough” in a given a situation. Simon referred to this behaviour as “satisficing”, and the form of rationality it represents as “bounded rationality” (Simon 1957).

In complex and fast-moving environments, boundedly rational individuals and groups cannot gather and cognitively process all relevant information before making decisions. Nonetheless, these decisions must be made, even in the absence of full information. What grounds do individuals use to make such decisions? According to Perrow (1972), we construct subjective, simplified models of a situation based on our past experience, and generally apply solutions we have used before. In other words, in the absence of perfect information, most of us fall back on heuristics. Heuristics provide simplifications that we can use to make decisions in “real time”, in situations in which gaining a full understanding of all relevant factors is either impossible or would require exhaustive analysis. They have thus have long been considered an important element in the discipline of behavioural decision theory.

Heuristics can often be helpful in solving certain classes of problems, although since we rarely know all of the factors actually governing a given turn of events, they make no guarantees (Perkins, 1981). Heuristics rely on a rough form of probabilistic reasoning, and are frequently used in games such as chess or cards. Chess masters use heuristics—such as “dominate the centre of the board”—to find acceptable solutions to complex problems without having to search the entire landscape of possibilities. Heuristic algorithms have also long existed in the field of computer programming, and have more recently been developed through the use of genetic algorithms and neural networks (e.g. Gupta et al 2000).

Different schools of thought have emerged concerning the value of heuristics, and the concept has become somewhat controversial. More specifically, some scholars believe heuristics should be uncovered and avoided (e.g. Kahneman and Tversky 1973), while others

believe that heuristics should be honed and used more extensively (e.g. Gigerenzer and Todd 2000).

Heuristics Should be Avoided

The first school of thought is dominated by social and cognitive psychologists, who in the 1970s conducted extensive studies on heuristics. The inspiration for this body of research was Tversky and Kahnemann's finding that individuals in psychology experiments depart from making decisions that rational choice theory would presumably recommend far too often for the cause to be random error. Further research led them to identify a number of heuristic "biases" or cognitive "errors"— shortcuts humans take that block decision makers from "effective decisions", and often lead to fallacious conclusions and errors of judgement. These conclusions were set out in the seminal work of Kahneman and Tversky (1973), who developed three main categories of such cognitive heuristics as biases:

Representativeness: we make a judgement based on how similar an incident is to a category that already exists in our minds. Our desire to categorize events or individuals according to our stereotypes sometimes leads us to ignore or give less credence to a priori probabilities or "base rates" (i.e. actual frequencies in the population). Some examples of this are illusion of validity or of control, reasoning by analogy, or being highly sensitive to vivid and anecdotal information.

Availability: we make judgements based on how easily we can bring to mind examples of the phenomenon at hand—at how "available" it appears to be. We tend to believe that easily retrievable phenomena are more common than they really are. Yet while often we do recall frequently occurring events more than less-frequently occurring events, it is certainly not always the case. For example, events that have affected or impacted us personally will come more easily to mind—we are biased—and in these cases we are likely to overemphasize their importance in making our decision.

Adjustment and Anchoring: we tend to make estimates by adjusting upward or downward from an “anchor”, or starting point. Thus, starting points can bias our decision making process. Such biases have been observed in many numerical judgments, such as house appraisals or assessments of geographic distances or the possibility of nuclear war. This bias has also been observed by social psychologists in for example, over-attributing an individual’s behaviour in what may be an artificial context to his or her underlying attitudes and dispositions. Outcomes of this heuristic can be escalation of commitment, single outcome calculations (focusing on the most promising possible outcome), and resistance to change.

Tversky and Kahneman (1986) concluded that heuristics are fundamentally and systematically push decision making away from pure rationality. The implication of their efforts is that heuristic biases are highly imperfect strategies that should be uncovered, understood, and then counteracted.

Several organizational scholars picked up on the work of Kahnemann and Tversky, and looked for evidence of heuristics in the management realm. Corporate acquisition and divestment decision making processes were associated with certain heuristic biases; for instance, reasoning by analogy, escalating commitment, illusion of control, and single outcome calculation (Duhaime and Schwenk 1985). The potential impact of biases such as availability, hindsight, sampling misunderstandings, correlation and causality, and representativeness was also explored in the context of strategic planning exercises (Barnes 1984). As with the social psychological approach, this body of work presumes that heuristics cause “faulty” judgements and lead to poor decisions, or at the very least, to overconfidence in a decision-maker’s heuristic-driven decisions.

Heuristics Should be Embraced

Other scholars have critiqued the notion that heuristics should be avoided at all costs, largely because the realist assumptions of perfect information and complete rationality rarely--if ever-

-exist. Presumably, conditions of pure rationality would make the “correct” course of action quite clear and obvious, thus eliminating the need for managers to reflectively interpret their business environments. Optimal management decision-making would be a highly generalizable skill, akin to following a pre-given recipe.

Where any of the key assumptions of the rational decision-making model is relaxed, decision makers have only bounded rationality. Making boundedly rational decisions requires the ability to simplify a complex situation sufficiently well to take context-specific action in real time. Due to our limited cognitive resources, and faced with a large number of decisions to be made each day, we “use approximate methods to handle most tasks” (Simon 1990: 6). In such an environment, well-developed and effective heuristics might be a highly efficient and effective way of making decisions. Although they are sometimes wrong, heuristics might prove useful if they save time in completing an exhaustive examination and evaluation of all possible alternatives, and if they provide correct or partially correct judgements more often than not. Heuristics can thus represent efficient uses of our limited cognitive abilities (Harvey 1998).

The comparison of heuristics to biases in the social psychology literature may also be somewhat of an oversimplification. Biases are presumed to be relatively static and deeply ingrained, whereas it is possible that heuristics might be steadily improved over time. Gigerenzer and Todd (2000) advocate the development of “fast and frugal” heuristics that are simple, robust in the face of environmental change, and generalize well to new situations. They further propose that heuristics may be a useful means of performing at least three tasks: guiding the search for information, stopping searching, and making decisions. Rather than trying to make optimal decisions, boundedly rational decision makers “bet” on simple heuristics for search and stopping rules.

Finally, rational choice theory also includes the implicit epistemological assumption that there is one, knowable reality. Yet if reality is truly socially constructed in a given cultural and institutional context (Berger and Luckmann 1966), it is particularly problematic to determine

what a pure “rational choice theory” may in fact indicate is an appropriate solution to a given problem. Such thinking underlies the pragmatist critique of realist epistemology (Rorty 1981), which implies that it is not possible to stand outside reality and gather all the relevant information to make an objective decision. Thus, everything that is not purely rational may not be irrational (Gigerenzer 1997). The “correct” course of action is rarely clear, as the complexity of today’s business environment increases faster than our cognitive abilities. In such a context, it would appear that the need for heuristic-like devices is greater than ever.

Questions Raised by the Literature on Heuristics:

Although the literature on heuristics is extensive, it leaves several questions unanswered:

How specific or general should heuristics be? The more specific heuristics are, the more of them are presumably needed to account for the multitude of contexts in which they could be called upon. Some scholars have suggested that just a small number of high-level heuristics are useful, such as in the area of ethics and morals (Soule 2000). Brady (1987) proposed that when developing and interpreting rules, we tend to fall back on just two “metarules”: the principle of benefaction (maximizing utility) and the principle of membership (equality, universalizability). Similarly, it has been claimed that “fairness” is a decision heuristic that helps individuals decide whether the directives of people in authority are fair, and thus, should be followed (Lind et al 1993). However, such heuristics may prove to be too general to provide significant guidance in particular situations--a frequently encountered problem with many statements of corporate values.

Should heuristics be solely internally focused, or externally focused, or both? Some debate appears in the literature concerning whether the usefulness of heuristics is more closely linked to internally or externally focused criteria. That is, should heuristics be judged based on their internal consistency and coherence, or on their usefulness in a given environmental context? While the inherent biases described by Tversky and Kahneman are necessarily

internally driven and focused, Gigerenzer and Todd (2000) favour externalist "ecological rationality" criteria, i.e. that heuristics should be tailored to the external environment, since what works in making decisions in one domain may not work in another.

Can heuristics be effective if imposed on others, or are they most useful if developed in a dialogical manner? In a given context, a heuristic can provide us with the rationale for making a decision, the "why" that is needed to justify a particular choice. While the social psychology literature implies that static, heuristic biases simply exist in our minds and should be uncovered and resisted, is it possible that heuristics can also be intentionally developed, or even imposed, on others? If so, how should their development process occur? To what extent can knowledgeable individuals and groups use heuristics to exploit the ignorance of others (Over 2000)?

Do heuristics exist purely in a cognitive, rational sense, or can they also be grounded in emotions? While most of social psychology literature deals almost exclusively with rational, cognitive heuristics, some interest has been generated by the notion of "hot cognition", that is, the notion that affect can influence our judgements (e.g. Forgas 1995). For example, falling in love has been proposed as part of a "stopping information search" heuristic (Gigerenzer and Todd 2000). Damasio has suggested that emotions play a biologically extensive and complex role in decision making, as our feelings mark some response options to real or simulated decisions, while unmarked options are not considered (Damasio 1994) However, little research has been done on emotionally-grounded heuristics. What role could emotional factors play in developing effective and robust organizational heuristics?

Method: Interpretive Case Studies

It is the purpose of this paper to shed light on organizational heuristics, and the four questions emerging from the literature review serve as a starting point for our exploration of this concept. However, we believe no study of organizational heuristics would be complete without

examining how they are actually used in the field. Thus, we went beyond the literature to examine how decision makers in organizations appeared to use heuristics in practice, by developing two interpretive case studies.

The first case study was LEGO Mindstorms[2], an autonomous, strategic project unit launched by an established toy company to develop a new business opportunity in an area on the boundary of the toy, Internet, and computer industries. Acting as participant-observers, we studied the team for a period of 22 months, from shortly after the team's establishment until the aftermath of the highly successful launch of the new product. Between May 1997 and February 1999 we conducted 27 in-depth interviews in five rounds of both members of the Mindstorms team and people in the parent LEGO organization. The interviews were semi-structured, following a prearranged set of headings and prompts but structure was otherwise kept to a minimum. The observation field notes were generated during discussions between sub-groups of the Mindstorms team and the researchers. Extensive secondary data was also collected and analyzed, including memos, press releases, drafts of organization charts, internal documents, and other reified organizational artifacts generated by team members during the period of study.

The second case study was AI Inc.[3], a New York-based company developing artificial intelligence applications for business. A total of six researchers observed and interviewed members of the AI Inc. team for 24 months from August 1999 to August 2001. We conducted 11 in-depth interviews with members of the organization, and also analysed e-mail data, internal memos, accounts prepared by team members for public consumption, drafts of strategic plans, and other internal and external documents.

We present the overview of each of the two cases, then identify what appeared to be the main heuristics that guided the interpretations and decisions of members of each organization. Our analysis makes pragmatic use of grounded theory to gain a better understanding of how these groups used heuristics to aid in decision-making in the context of organizational action.

[2] The names of individuals involved in the LEGO Mindstorms case study have been disguised.
[3] All company, individual, and product names in the "AI Inc." case study have been disguised.

While there has been some debate concerning whether or not decision-making and action are separate activities following separate rationalities (Brunsson 1982; Reed 1991), we believe they represent two sides of the same coin: action presupposes decision-making and decision-making generally leads to action.[4]

We use these findings, together with the results of the literature review, to extract plausible relationships among concepts and sets of concepts (Strauss and Corbin 1998:168). Our findings from the case studies shed light on the four questions that emerged from the literature review, and facilitated our generation of four empirically visible dimensions of heuristics—scope, focus development and depth. The case studies also lead to the development of four propositions that describe how organizational heuristics might be effectively used in an organizational setting. The research process was thus both inductive and deductive; the two interpretive case studies informed the theoretical reading of the literature on heuristics, and the findings of the literature review informed the empirical work.

Because the research approach is interpretive, traditional notions of internal validity, construct validity, reliability and replicability, and external validity (Yin 1994) are problematic. Reliability and replicability in particular require the existence of invariant laws over time (Numagami 1998). In their place, a number of other criteria have been proposed with which to evaluate interpretive research in the social sciences. Taylor (1971) asserted that we need to move beyond a science of verification to one that studies how intersubjective and common meanings are embedded in social reality. It is broadly agreed that in rejecting traditional validity tests assuming objectivity, the credibility of the researcher takes on much greater importance for validity in interpretive case study development. Denzin and Lincoln (1998) emphasize the importance of the researcher, proposing that the trustworthiness and authenticity of the work should replace traditional positivist criteria of internal and external validity.

[4] In a sense, two levels of heuristics were at play: heuristics in use within the organizations, and the heuristics we as researchers used to identify their organizational heuristics. In identifying heuristics, we focused on commonly-cited criteria used in sensemaking for action or decision-making.

***Case Study 1*: LEGO Mindstorms**

In 2001, the LEGO[5] company was the largest toy company in Europe, and the fourth largest in the world. From its founding in 1932 in the village of Billund, Denmark, LEGO had become one of the best-known brands on the planet, with fully 64% of US households and 74% of European households owned LEGO toys. Throughout the 1980s and early 1990s, LEGO had grown at a controlled pace of between 10-18% annually, although the rate of sales growth began to slow in the late-1990s. At this time, the company was facing a growing threat from new competitors and the challenge represented by the rapid spread of the Internet and technology-oriented toys. LEGO management had found it difficult to encourage the development of radically different products to counter these challenges.

In April 1996, Sam Sinclair left his job with an information technology company to take over the project and found a new team with LEGO product manager Frank Fulton. The Mindstorms team was set up as an autonomous project unit, reporting directly to the senior management of LEGO, but located within part of one of the least profitable businesses of the company, far from the view of most of the organization. “We started in the corner of the corner”, recalled Fulton. The project was considered important because it would potentially represent a new direction for the parent company, taking it into what was at the time the relatively unexplored realm of computerized toys. Yet it was expected to be difficult, as the company had little previous experience in dealing with computer-based products.

The two project founders were given a mandate to take some partially finished technology, develop an original product, and bring it to market. From the beginning, they set out to operate in a way quite different from the parent company. The team operated using highly unconventional procedures, as they claimed to disregard hierarchy. “We threw away most of the rulebook, and operated like a small, entrepreneurial business,” said Fulton, distinguishing the team from its parent company.

[5] The name “LEGO” was chosen as a contraction of two Danish words: “leg godt” meaning “play well”.

The centerpiece of LEGO MINDSTORMS would be the “RCX” programmable brick, which used sensors to take in information from its environment and signal output motors to turn on and off. Children could use the RCX along with other LEGO pieces to build robots that could be programmed to follow lines, clear a table, or perform almost any task. Kids would also be connected in a “global community”, focused around an Internet website: www.legomindstorms.com. The MINDSTORMS team planned to largely forgo traditional distribution channels in favor of consumer electronics stores, special chains for learning, and directs sales through LEGO clubs, LEGO MINDSTORMS learning centers, and the Internet. The team also faced an unfamiliar lineup of competitors; instead of Mattel and Hasbro, MINDSTORMS would compete with products from Microsoft, Sony, Nintendo, and Sega.

Over the two-year period, the team grew from its original two members to include a total of 15 individuals representing several nationalities from a variety of different parts of the organization, industry backgrounds, and locations in three different countries. Many of these people came from outside LEGO. Sinclair quickly ran into problems when he tried to recruit employees from within the LEGO organization to work for the fledgling team. He recalled “Many people didn’t want to be too closely associated with MINDSTORMS…they wanted to maintain some distance in case it failed”. It took a great deal of effort for the team to finally acquire a project manager and public relations specialist from the core business. In many cases Sinclair had to resort to hiring external people rather than draw on the knowledge within the LEGO organization. In February 1997, he hired a US-based Internet specialist to develop the project’s Internet community, and a few months later hired a marketing director from an educational software company, who would also be based in the US.

In January 1998, LEGO MINDSTORMS conducted its global public relations launch. The event was covered by one-half of all US television stations, as well as CNN and the BBC, and was estimated to have reached more than one billion people around the globe. In September 1998 the product went on sale in the United States and the UK, and sales doubled projections—stores were sold-out by Christmas.

Early on, the MINDSTORMS team set out some operating principles that contrasted sharply with those of the core LEGO organization. The team would meet regularly and make extensive use of external alliances—not just internal ones—in order to keep the center “lean”. It also set out to treat its alliance partners—many of which were far larger than LEGO Mindstorms—as equals, working with them collaboratively rather than through a relationship of domination. They would strive to run processes in parallel—and across boundaries—instead of sequentially, and allowed more “chaos” than was permitted in the core business. For example, the team worked to develop the hardware and software for the product simultaneously, a style of working that would never have occurred in the parent company.

An enthusiastic “can-do” attitude infused the fast-moving team, which knew it had to achieve results quickly in order to survive in and create its new category in an emerging industry. But the fast pace of the team sometimes took a toll on its members. Recalled Askildsen: “I was getting exhausted from all our parallel processes—it never stopped. The excitement was energizing but it can wear you out.”

A common subject of debate was how MINDSTORMS should be connected to the rest of the organization. In August 1997, after spending a great deal of time and effort building contacts with managers in the core business, Sinclair designed a highly elaborate organization chart showing how his team “fit in” with the rest of the LEGO organization. However, the boxes and arrows of an organization chart seemed inadequate to describe the fundamental differences between the orderly core business that was servicing an existing market on the one hand, and the newer more chaotic team trying to create a new market on the other. Majgaard began adopting a more metaphorical way of describing the linkage, referring to MINDSTORMS as a “planet” orbiting around the core business. He emphasized that it was important for all new business development ventures to “stay in orbit”, not drifting too far away into oblivion, nor getting too close and burning up. As could be expected, this balance was not always easy to maintain.

Heuristics in LEGO Mindstorms

Although their environment was highly uncertain, and at no time did the Mindstorms team have all the possible information available with which to make decisions, it did take made them on a regular basis in real time. We list below the heuristics to which members of the Mindstorms team made reference in describing their decisions and actions:

Stay in orbit: the use of this metaphor by all team members—as well as the parent company—enabled the team to make sense of its relationship with Lego: linked by an invisible yet powerful gravitational force. This principle came into play when the team was developing a common public relations statement with the parent company, while striving to maintain its own identity. It also helped the team better make sense of how it could draw on the well known Toys name while not having to adopt the parent company's operating practices.

Parallel process: the team explicitly set out to avoid the step-by-step, "follow the game plan" approach to operating of its parent company, in favour of working on several tasks at the same time. In this way, it reflected the constantly changing business environment it was facing. An engineer explained this heuristic increased the team's speed: "because most of the development progress has been done in parallel, we haven't had to wait for each other to finish things before we could continue with others".

Be a real partner: the team did not feel the need to assume a dominant position with their partners, which allowed it to partner with a wider variety of both large and fledgling partners. Working with partners in this way is a critically important starting point for building a real network. Even beyond the completion of the product, the team knew that sales could be boosted if they could create a network of customers, which in turn could be facilitated by creating a large network of partners.

Carpe diem: the team was action oriented at all times. By eschewing traditional planning processes, the team was able to make decisions more quickly than its parent company. Members of the team would proactively take initiatives in their own domains, including establishing new partnerships and dealing with existing ones, rather than waiting for approval from the team leader.

***Case Study 2*: AI Inc.**

Formed in 1997 through the collaboration between artificial intelligence scientist Bob Goodwin and Wall Street economist Laura Pierce, AI Inc was established to build the world's first thinking machine. Although artificial intelligence had attracted much interest and investment in the 1980s, it was an elusive goal, and AI 's ability to deliver on its initial high expectations meant it had been largely ignored by the investment community. Goodwin, however, had developed what he claimed to be a radically different approach to AI based on an evolving series of "nodes" and "links" representing concepts, which were designed to emulate the processes of the human mind. It's natural language capabilities allowed this new AI system—called "Wow"—to understand what information *meant* in the context of the user and the enterprise. In theory, it could thus tap into the estimated 90% of the world's stored business information residing in text documents, e-mails, and so on—information that previously could only be analysed as fast as decision-makers could read it.

Goodwin began focusing intently on developing the Wow architecture, which he anticipated could be completed within just a couple of years. However, he and Pierce soon realized that in order to finance the development of Wow, they would have to spin off some applications or artificial intelligence "components" along the way. For example, they had worked on an intelligent web search engine that would conduct meaning—instead of keyword—searches. However, early tests of the search application showed it didn't live up to expectations, so work began on another application that would predict movements in the financial markets. Under

the assumption that news drives markets, this Wow application would read financial data off the Web, extract and interpret key concepts, and look for patterns and clues with which to make predictions and trading recommendations. In 1998, Pierce formulated a strategic plan aiming the market prediction application of Wow squarely at potential customers in the financial industry, with the intention of finding a market where the technology would be expected to generate immediate profits.

This plan helped attract the interest of Arthur Smith, the former Head of Fixed Income and Treasuries at a major European bank, who would go on to take over from Pierce as AI Inc's CEO in March 1999, Smith, who had developed long-standing relationships with several venture capitalists, quickly arranged another round of venture capital funding, which allowed a wave of hiring of computer programmers to speed progress towards the completion of Wow. Smith believed AI Inc. should move rapidly into the Internet space, and began exploring possibilities of getting co-development funding with clients who could then use the technology in specific industry domains. However, it was a challenge for the company to "close the gap of credibility" with potential partners, so he developed "valuation events" at which the technology was exposed to potential investors. The need to develop such application demos put a strain on Goodwin, who placed relatively low priority on such tasks. Goodwin explained:

I need to be involved in getting the thinking machine to work, plus I need to do the dog and pony show for investors and potential customers, and so managing application development doesn't get as much attention as it should.

The situation became more complicated when the team was approached by Matthew Lightfoot, a former Wall Street investment banker, to build a Wow-enhanced knowledge management application for legal firms. The company's typical reaction to these new opportunities was to "reorient and redirect" the medium term strategy as they arose, a strategy that led to considerable frustration among the product developers. While Pierce saw its goal as "to become a very successful corporation, successful meaning very high valuations and huge revenues", Goodwin saw the company's overall goal as building the world's first

thinking machine, believing that the finishing the artificial intelligence system should come ahead of developing saleable components. This difference led to fundamental differences of opinion concerning how to allocate resources. Pierce and Smith prioritised making new deals for applications, believing that instead of caring about the technology inside, customers cared about how it could solve a pre-existing problem. On the other hand, Goodwin believe that they should focus on the long-term goal and just work towards it, believing that "if we build a better way of predicting the markets, the world will beat a path to our door".

Financially, AI Inc. began with an initial investment from founders, who attracted an initial round of seed funding in March 1998. An article in a prominent newspaper helped them generate a second round of funding in the middle of 1998, although the market downturn in late 1998 meant they almost ran out of money. This prompted a renewed focus on developing saleable applications in 1999, and they were able to generate more funds based on huge market valuations placed on similar companies at the time. However, the company avoided raising too much money, as it was concerned that raising too much money from venture capitalists would dilute the founders' share. Yet expenses were increasing rapidly.

As the company grew from an initial group of five founders to reach 130 employees at its peak in 2000, the recruiting style of the organization appeared to change. In the early days, Pierce characterized their approach as "taking inexperienced people we thought were really smart and hire them in for nothing". Goodwin called it "looking for the most brilliant people with the most ideas of their own which relate to what I do", and sought out to attract all kinds of fascinating, brilliant, independent minds to work on Wow. The idea was to give individuals the maximum control over their own lives, and provide minimal hands-on management. However, the two founders realized over time that they needed individuals who, while less brilliant, had more specific experience. As Pierce described: "everyone doesn't have to be like me".

Decision Making at AI Inc.

Both Pierce and Goodwin had what was characterized as having a “ready, shoot, aim” approach to making decisions; trusting their intellects and moving quickly. Unlike most technology start-ups in which the creator of the technology hands over its development to other managers and moves into a research role, Goodwin carried much of Wow’s architecture in his head, and maintained considerable control over its development of Wow. Pierce also wanted to be involved in everything, however as the business of running AI Inc—and the technology itself—became more complex, this “mom and pop” approach to running the business became more difficult. The founders’ opinions often contrasted with those of some of the newer, more “professional” managers they brought in, who saw more value in developing plans.

The speed with which AI Inc. had grown meant that it had trouble finding enough highly qualified employees for its New York office, and began seeking to employ people overseas. Visa restrictions for the US made it practical for these new hires to remain in their home countries and conduct their work and a large part of their communication with AI Inc. over the Internet. Over time, AI Inc.’s “distributed” organization structure meant it had employees in Brazil, Australia, New Zealand, South Africa, and Russia. Most of these employees had seen Goodwin’s website and contacted him directly; many had never set foot in AI Inc.’s offices, or even met one of the company’s founders. Many worked on specific tasks in isolated locations. Even Goodwin, Pierce, and Smith only came into the office occasionally, mostly they worked from home.

The distributed structure meant that e-mail was an important communication medium for employees, with Goodwin receiving an average of 400 per day. As one staff member remarked: “We used email to communicate, which is good for documenting things, but we sometimes miss secondary information…many of us come from academia, and we are mimicking the communication method common there”. To have some face-to-face discussion time, the company decided to meet in the office every Tuesday and Thursday (Tuesday for

technology questions, Thursday for management discussions). Despite the fairly “constant chatting”, arranging meetings was difficult. As one senior manager described: “we talk about getting the executive committee together on a formal basis to discuss stuff, but somehow it always gets bogged down in a million different things, so we don’t do it.” As the office manager recalled: “scheduling four people for a meeting can be a half-day job”.

In early 2000, the company hired an external consulting company to help it generate a strategic plan, and focus its efforts. The plan spelled out that AI Inc. would become a “component specialist”, focusing on developing an intelligent search engine and selling it to intermediaries. The company believed that this strategy would allow it to continue with its research while giving up the minimal amount of control by selling products to finance its continued operations. However, the bursting of the dot.com bubble would deal AI Inc. a fatal blow. Their partially developed component strategy meant that by the time they were almost ready to sell a particular component such as intelligent web search, most of the Internet companies who would have been potential buyers had gone bankrupt. The company had placed great hopes in negotiating a deal with a major European media company, but the deal fell through, leaving AI Inc. to engage in round after painful round of layoffs, and ultimately declaring bankruptcy in April 2001.

Heuristics Used at AI Inc.

Hire based on intelligence. Rather than basing their hiring decisions on how well the candidate fit into the existing organization, Pierce and Goodwin believed in hiring individuals based on their raw talent. Pierce called it “taking inexperienced people we thought were really smart and hire them in for nothing”. Goodwin called it: "look for the most brilliant people with the most ideas of their own which relate to what I do”. Later Pierce tried to revise this, but by then it was too late.

Avoid dilution. An important heuristic adopted by the founders of AI Inc. early on was their clear desire to avoid giving up control of the company (dilution) at all costs. At several points in interviews with the company's founders, it was clear that many potentially beneficial alliances were being refused based on this heuristic, which is evident in the following excerpt of the company's April 2000 strategic plan: "This option (securing a large investment to conduct 2-5 years of R&D) was deemed undesirable because the mount of funds required to pursue it would require us to immediately hand over control of the company to the major investor". At another point, Goodwin remarked: ""We never had a lot of extra money in the bank, but we always had the confidence we could raise more quickly, if necessary".

Meet together twice per week. AI Inc was a distributed organization, meaning that at any one time relatively few of its total employees would be co-present in the company's' New York office. Even the management committee tended to only come into the office occasionally, with Goodwin and Pierce generally working from their homes in neighbouring states, making it difficult to arrange meetings. The vast majority of communication within the company occurred through e-mail. Thus the team made the decision to be present in the office two days per week, Tuesdays for technology questions, Thursdays for management decisions.

Build the AI machinery first. This heuristic was evident in statements such as "after Wow is finished, the rewards will take care of themselves and we will be richly compensated in a variety of ways". It was important in that it prioritised finishing the core AI product before worrying about business applications. This also meant the company did not spend much time seeking out customers for its product. As Goodwin stated: *"…*until someone writes me a cheque, let's know what our long term goal is and just work toward it". This heuristic was mainly espoused by Goodwin, whereas Pierce and the business side wanted to focus more on developing business applications first, to pay the bills.

Discussion and Propositions

Each of the two organizations we studied employed numerous heuristics during the study period. In identifying what we interpreted to be the four main heuristics used in each organization, we observed that these heuristics varied along at least four "heuristic dimensions": scope, focus, development, and depth (*see Figure 1).* We now turn to discuss each of these dimensions.

Figure 1: Dimensions of Identified Heuristics

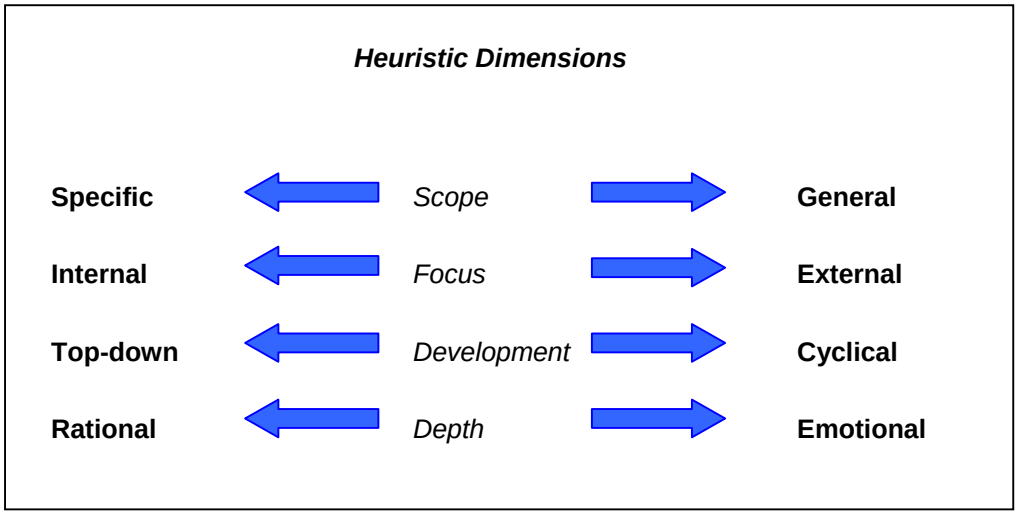

First Heuristic Dimension: Specific-General

The heuristics we identified appeared to vary in terms of their specificity. Highly specific heuristics were those that had a very precise meaning and provided clear direction on what needed to be done, while more general ones were more distant from specific decisions, and required a greater degree of interpretation. Of the eight heuristics we identified in the two cases, two appeared quite specific ("avoid dilution" and "meet together twice per week"), while the other six—including all four LEGO Mindstorms heuristics—appeared to be quite general (for example, "stay in orbit").

The idea of managing based on a few specific rules has been explored to some extent in the complex adaptive systems literature under the notion of "flocking" [6]. In a computer simulation, Reynolds (1987) discovered that of complex global behaviour emerged based on the interaction of a few highly specific rules. Although his computer-simulated birds or "boids" were not programmed to flock, they did so, even though it would have been impossible to predict such an outcome merely by looking at the three rules. The important message from this work is that flocking occurred even though no individual boid only saw what was happening in its local region--none had an overall view of what was happening. Nor did they formally co-ordinate their activities or follow the instructions of a leader.

While the notion of emergence from simple rules at first prompted great interest among management scholars, these simplified applications of complex adaptive systems theory have been criticized by some social scientists, who make the point that human beings are far more complex than simple boids. We rarely if ever could base our actions solely upon three specific rules. Nonetheless, some management scholars have picked up on the concept of simple rules. Eisenhardt (2001) appears to build on these studies in developing her recommendations that firms should develop heuristics resembling "simple rules". She claims that advantage comes from seizing fleeting, unexpected opportunities, and that managers should pick "a small number (from 2-7) of strategically significant processes and craft a few simple rules to guide them."

Narrowly defined and highly specific heuristics are less likely to be robust to a variety of environments, and may lead an organization to adopt overly rigid and predictable responses in the face of changing business imperatives. Eisenhardt's statement that "it is important for companies with simple rules strategies to follow the rules religiously—think Ten Commandments, not optional suggestions—to avoid the temptation to change them too frequently" (Eisenhardt 2001:115), is somewhat contradicted by her later comment that they sometimes have to be changed if the main strategy "goes stale". Highly specific heuristics may need to be changed more frequently if they are inflexible to an organization's changing

[6] It has been found that computer simulations of birds "boids", will flock consistently when programmed with just three key rules: keep a minimum distance from other objects, try to fly at the same speed, and head where the concentration is densest in the vicinity.

context and identity. In addition, competitor firms may find it easier to copy very specific heuristics. Heuristics cannot be reduced to a local or isolated problem because they are about how we should interpret a situation (Kessels 2001).

On the other hand, more general heuristics are designed to be robust to a variety of contexts. General heuristics may be expected to guide behaviour while allowing for the emergence of entirely new phenomena. They are perhaps better able to cope with the loosely coupled linkages between decision and action, changing external factors, and the impact of political and symbolic considerations that drive most decision-making in practice (Reed 1991). March (1989) has claimed that decision-making is a highly contextual, sacred activity, surrounded by myth and ritual, and as much concerned with the interpretative order as with the specifics of particular choices. In complex and fast-changing environments, decisions may require more of a "bricolage" approach of interpreting signs, rather than "engineering" process for dealing with transparent concepts (Levi-Strauss 1966). In such environments, heuristics need to be general enough to allow for interpretation.

Proposition 1: Because they allow for interpretation, general heuristics will be more robust and effective than specific heuristics in complex business environments.

Second Heuristic Dimension: Internal-External

The heuristics we identified also appeared to vary in terms of their focus. Several appeared to deal mainly with the internal workings of the group, while others appeared to focus on the environment in which the group operated. Of the eight identified, only two ("be a real partner" and "stay in orbit") appeared to address the environment, while the other six—including all of the AI Inc. heuristics—appeared to focus on the internal workings of the group.

On the individual or group level, our heuristics are in part socially constructed, and change through our interactions with other. For example, we may be influenced in ethical questions

by factors in our environments, our friends, families, legal systems, team members, and so on. Thus, rationality that is purely self-referential may not lead to effective decisions in complex environments. Gigenrenzer and Todd (2000) emphasize the importance of context-dependent ecological rationality in determining the usefulness of heuristics. Group rationality that ignores the task environment runs the risk of becoming an elaborate exercise in navel-gazing.

Heuristics that embrace the outer environment can also send a clear and signal to partners and other parties in the team's environment. For example, to the extent that the Mindstorms principle of "work as equal partners" operated as an informal psychological contract, it may have guided the team towards a way of working that may have required fewer formal contractual safeguards between the team and its partners (Ring and Van de Ven 1994).

Proposition 2: A balance between internally and externally focused heuristics will improve a group's decision-making processes by making them more contextually rational.

Third Heuristic Dimension: Top-down—Cyclical

The heuristics we identified differed in the nature of their development. Interestingly, all of the AI Inc. heuristics were developed in a top-down manner, with the senior management group imposing the heuristics on the organization. In some cases, there was disagreement among this group on the heuristic, and occasional conflict was evident, for example, in the "build the AI machinery first" heuristic. On the other hand, the Mindstorms heuristics were developed in the course of discussions among team leaders and members through a "cyclical" process. For example, "stay in orbit" was articulated time and time again by various members of the team, and was frequently used to help clarify the team's relationship to the parent company, and thus, how it should act. We believe that heuristics developed in such a cyclical, dialogical

manner are more likely to lead to a sense of "ownership" (Reason and Rowan 1981) of the heuristic by all individuals concerned.

Dating back to Weber's bureaucratic iron cage, there has been a persistent belief that instrumental techniques are more important than reflective discussion about organizational values in organizational decision-making (Reed 1991). We argue that such thinking tends to be grounded in rational analysis, and leads to rule-like heuristics that are typically imposed by top management in organizations in much the same manner as suggested by Eisenhardt (2000). In this vein, we note Simon's (1993) belief that, in order to be effective, mission statements must become part of the mind-set of every member of the organization who is responsible for making or helping to make decisions of any consequence.

However, heuristics that emerge in a cyclical manner in teams can also be tightly enforced (Barker, 1993), and they are more likely to tap into the tacit knowledge (Polanyi 1966) of team members than those imposed top-down. Such heuristics are also more likely to encourage both participation *and* reification (Wenger 1999), thus improving the quality of engagement of team members. Such heuristics are seen more as processes than reified objects. When leaders carefully consider team member inputs, the members see the process as fairer, have greater commitment to the decision, greater attachment to the team and greater trust in the leader (Korsgaard et al 1995). However, moving from rules to principles requires an element of trust. Organizations with excessive top-down imposed rules and regulations are typically characterized by trust disorders and paranoid thinking (Kets de Vries 1999). It appears that the Mindstorms team was very effectively able to develop heuristics through a process that was both top-down and bottom-up.

Proposition 3: Heuristics should be developed in a cyclical manner through regular dialogue of group members.

Fourth Dimension: Rational—Emotional

Finally, the heuristics we identified differed in the extent to which they were driven by rational, as opposed to emotional, analysis. Based on our interpretation of our data collection, we believe that three of the heuristics ("be a real partner", "stay in orbit" and "carpe diem") were grounded in emotion, while the other five—including all of the AI Inc. heuristics—appeared grounded in rational analysis. We refer to this dimension as "depth" under the assumption that heuristics driven by cold analysis are less likely to be internalised and used than those driven by hot—or at least warm—analysis.

The sharp distinction we make between reason and emotion dates back to Plato. However, it increasingly appears that emotions and rational thinking are interlinked. For example, emotions appear to play a significant role in the process of scientific discovery (Thagard 2001). Achieving positive affect within a team has been shown not to adverse influence the quality of decisions (Korsgaard et al 1995). Yet most literature on emotions and management studies imply that emotions are something to be managed and suppressed, principally on the grounds that they interfere with rationality. Other scholars have proposed that emotions can serve rationality, suggesting for example that negative emotions such as anger, serve to ensure the credibility of threats, while love and affection ensure the credibility of promises to cooperate (Fineman 1999:296). However, while some managers have attempted to control the emotions of subordinates, for example—through customer service training initiatives, such efforts can be problematic in that the emotions elicited are not genuine. These attempts may also signal a return to mechanistic models of management. "Deliberate hierarchical control of emotions—where positive, "smiling" ones are inculcated and negative 'stress' ones are pathologized—presents a compelling picture of a late twentieth century form of Taylorism." (Fineman 1999:303).

Emotional attachment to heuristics may develop if the heuristic is consistent with an individual's and/or a group's *values*. Such a notion appears to align with Lissack and Roos' (1999) notion of "simple guiding principles", which are ideally grounded in values, and aligned

with individual and organizational identities. These authors recommend that firms explicitly develop such heuristics in order to improve organizational coherence. It would appear that heuristics grounded in values and identity may generate a stronger emotional attachment than those that are not. Thus, they are more likely to be internalised and made effective by group members.

Proposition 4: ***To be most effective, heuristics should not be purely rational; they should also take account of our emotional selves.***

Conclusion and Further Research

In a complex, ever-changing business environment in which it is impossible to always know everything that is going on, truly optimal decisions are extremely difficult to make. This fact underlies the notion of inferential heuristics (Kahneman and Tversky 1973), rules of thumb we all use to make decisions in our daily lives. Organizations and teams that develop meaningful heuristics that can help people make the many important decisions they face each day will be able to guide behaviour more dynamically and robustly than through use of a series of static rules. Effectively used, a set of shared guiding principles can help facilitate the development of a robust “collective mind”.

However, this work leaves many questions unanswered. How do guiding principles emerge? How can they be intentionally developed? Are guiding principles more of a sensemaking tool or do they provide performative benefits for a group (i.e. direct aid in decision making)? Are guiding principles really only meaningful to the individuals involved in their development? We believe further interpretative research may help shed light on some of these many remaining questions in the field of organizational heuristics.

REFERENCES

Allison, G.
1971 *Essence of decision.* Boston: Little, Brown and Company.

Barnes, J.
1984 'Cognitive biases and their impact on strategic planning'. *Strategic Management Journal* 5/2: 129-137.

Barr, P., Stimpert, J. and Huff, A.
1992 'Cognitive change, strategic action, and organizational renewal'. *Strategic Management Journal* 13: 15-36.

Berger, P. and Luckmann, T.
1966 *The social construction of reality.* New York, Doubleday.

Brady, F.
1987 'Rules for making exceptions to rules'. *Academy of Management Review* 12/3: 436-444.

Brunsson, N.
1982 'The irrationality of action and action rationality: Decision, ideologies and organizational actions'. *Journal of Management Studies* 19/1: 29-44.

Damasio, A.
1994 *Descartes' error: Emotion, reason and the human brain*. New York: Grosset/Putnam Book.

Denzin, N. and Lincoln, Y.
1998 'Major paradigms and perspectives' in Denzin, N. and Lincoln, Y.(eds), *The landscape of qualitative inquiry,* 185-193, London: Sage.

Duhaime, I. and Schwenk, C.
1985 'Conjectures on cognitive simplification in acquisition and divestment decision-making'. *Academy of Management Review* 10/2: 287-295.

Eisenhardt, K.
2001 'Strategy as simple rules'. *Harvard Business Review*. 79/1: 107-116.

Fineman, S.
1999 'Emotions and organizing' in Clegg, S. and Hardy, C. (eds) *Studying organization: Theory and Method*, 289-310. London: Sage.

Forgas, J.
1995 'Mood and judgement: the affect infusion model (AIM)'. *Psychological Bulletin*, 117: 39-66.

Gigerenzer, G., and Todd, P.
2000 *Simple heuristics that make us smart*. New York: Oxford University Press.

Gigerenzer, G.
1997 'Bounded rationality: Models of fast and frugal inference'. *Swiss Journal of Economics and Statistics,* 133/2: 201-218.

Glaser, B. and Strauss, A.
1967 *The discovery of grounded theory.* Chicago: Aldine.

Gupta, J., Sexton, R. and Tunc, E.

2000 'Selected scheduling heuristics using neural networks'. *INFORMS Journal on Computing* 12/2: 150-162.

Harvey, J.
1998 'Heuristic judgment theory'. *Journal of Economic Issues* 32/1: 47-64.

Kahneman, D. and Tversky, A.
1973 'On the psychology of prediction'. *Psychological Review* 80: 237-251.

Kessels, J.
2001 'Socrates comes to market'. *Reason in Practice*, 1/1: 49-71.

Kets de Vries, M.
1999 'High performance teams: Lessons from the Pygmies'. *Organizational Dynamics* 27/6: 66-77.

Korsgaard, M, Schweiger, D. and Sapienza, H.
1995 'Building commitment, attachment and trust in strategic decisions'. *Academy of Management Journal* 38/1: 60-84.

Lévi-Strauss, C.
1966 *The savage mind*. Chicago: University of Chicago Press.

Lind, E., Kulik, C., Ambrose, M., and de Vera Park, M.
1993 'Individual and corporate dispute resolution: Using procedural fairness as a decision heuristic'. *Administrative Science Quarterly* 38/2: 224-251.

Lissack M. and Roos, J.
1999 *The next common sense.* London: Nicolas Brealey.

March, J.
1989 *Decisions and organizations*. Oxford: Basil Blackwell.

Oliver, D. and Roos, J.
2001 'Coherence in collaboration: The case of self-managed teams'. *Imagination Lab Working Paper*, WP2001-3.

Over, D.
2000 'Book review: Ecological rationality and its heuristics'. *Thinking and Reasoning*, 6/2: 182-192.

Perkins, DN
1981 *The mind's best work.* Cambridge, MA: Harvard University Press

Perrow, C.
1972 *Complex organizations*. New York: Scott, Foresman.

Polanyi, M.
1966 *The tacit dimension*. London: Routledge.

Pruzan, P.
2001 'The question of organizational consciousness: Can organizations have values, virtues and visions?' *Journal of Business Ethics* 29/3: 271-284.

Reason, P. and Rowan, J. (eds)
1981 *Human Inquiry: A Sourcebook of New Paradigm Research.* Chichester: John Wiley.

Reed, M.
1991 'Organizations and rationality: The odd couple?' *Journal of Management Studies*, 28/5: 559-567.

Reynolds, C.
1987 'Flocks, herds, and schools: A distributed behavioral model'. *Computer Graphics* 21/4: 25-34.

Ring, P. and Van de Ven, A.
1994 'Developmental processes of cooperative interorganizational relationships'. *Academy of Management Review* 19/1: 90-118.

Rorty, R.
1981 *Philosophy and the mirror of nature*. Princeton, NJ: Princeton University Press.

Simon, H.
1957 *Models of man: Social and rational*. New York: Wiley.

Simon, H.
1993 'Strategy and organizational evolution'. *Strategic Management Journal* 14: 131-142.

Soule, E.
2002 'Managerial moral strategies: In search of a few good principles'. *Academy of Management Review* 27/1: 114-124.

Stanwick, P.
1996 'Mental imagery: An alternative to top management team replacement for declining organizations'.*Journal of Organizational Change Management* 9/2: 47-65.

Strauss, A. and Corbin. J.
1998 'Grounded theory methodology: An overview' in Denzin, N. and Lincoln, Y. (eds) *Strategies of qualitiative inquiry,* 158-183. London: Sage.

Thagard, P.
2001 'The passionate scientist: Emotion in scientific cognition'. (forthcoming: available at: http://cogsci.uwaterloo.ca/Articles/Pages/passionate.html)

Todd, P.
2000 'The ecological rationality of mechanisms evolved to make up minds'. *American Behavioral Scientist* 43/6: 940-956.

Tversky, A. and Kahneman, D.
1986 'Rational choice and the framing of decisions'. *Journal of Business*, 59/4: S251-S278.

Wenger, E.
1999 *Communities of Practice: Learning, Meaning, and Identity.* Boston, MA: Cambridge University Press.