

Empirical tests of a fast-and-frugal heuristic: Not everyone “takes-the-best”

Ben R. Newell,* Nicola J. Weston, and David R. Shanks

Department of Psychology, University College London, Gower St., London WC1E 6BT, UK

Abstract

The fast-and-frugal heuristics approach to decision making under uncertainty advocated by Gigerenzer and colleagues (e.g., Gigerenzer & Goldstein, 1996) has achieved great popularity despite a relative lack of empirical validation. We report two experiments that examine the use of one particular heuristic—“take-the-best” (TTB). In both experiments the majority of participants adopted frugal strategies, but only one-third (33%) behaved in a manner completely consistent with TTB’s search, stopping and decision rules. Furthermore, a significant proportion of participants in both experiments adopted a non-frugal strategy in which they accumulated more information than was predicted by TTB’s stopping rule. The results provide an insight into the conditions under which different heuristics are used, and question the predictive power of the fast-and-frugal approach.
© 2002 Elsevier Science (USA). All rights reserved.

Keywords: Bounded rationality; Fast-and-frugal heuristics; Take-the-best; Behavioral strategies

1. Introduction

The ‘fast-and-frugal’ approach to judgment and decision making has achieved recent wide popularity (Gigerenzer, 2000; Gigerenzer & Goldstein, 1996; Gigerenzer & Selten, 2001; Gigerenzer & Todd, 1999). The approach has been applied to many areas including legal decision-making (Dhimi & Ayton, 2001), decision analysis in patient care (Elwyn, Edwards, Eccles, & Rovner, 2001), decision making in colonies of social insects (Seeley, 2001), and has also been tested in a number of studies employing more traditional cognitive psychological paradigms (Bröder, 2000; Jones, Juslin, Olsson, & Winman, 2000; Newell & Shanks, 2003).

The heuristics that are grouped under the fast-and-frugal banner are realizations of the probabilistic mental models (PMM) theory of human inference (Gigerenzer, Hoffrage, & Kleinbolting, 1991). One of the inspirations for PMM theory was Simon’s (1956) concept of “bounded rationality”—the term he used to highlight the interconnectedness of the limitations of the mind

and the structure of the environment in his model of human decision-making. Simon argued that people ‘satisfice’—or look for ‘good enough’ solutions that approximate the accuracy of optimal algorithms without placing too heavy a demand on the cognitive system (e.g., working memory, processing capacity). The major thrust of both Simon’s and subsequently Gigerenzer and colleagues’ approach is that the emphasis on speed and frugality replaces the methods of classical rationality (e.g., expected utility theory and Bayesian reasoning) with “simple, plausible psychological mechanisms of inference...that a mind can actually carry out under limited time and knowledge” (Gigerenzer & Goldstein, 1996).

This seductive appeal of heuristics that are both simple and plausible has undoubtedly contributed to their popularity. For example, Elwyn et al. (2001) have argued that there needs to be a paradigm shift in the way decisions are managed that reflects the importance of analyzing the trade-off between accuracy and frugality in the practice of evidence-based medicine. They argue that although methods of decision making based on classical approaches embodied by Bayesian reasoning and expected utility theory (e.g., decision-analysis) might improve patient autonomy by involving patients

* Corresponding author. Fax: +44-0207436-4276.

E-mail address: b.newell@ucl.ac.uk (B.R. Newell).

in the decision process, their use by health professionals with limited time is wholly unrealistic. Elwyn et al. (2001) propose that the ‘new frontier’ will involve the development of fast-and-frugal heuristics for both clinicians and patients to enhance the transparency of diagnosis and decision-making.

Is this talk of ‘new frontiers’ and paradigm shifts warranted? In this paper we are interested not so much in whether fast-and-frugal heuristics are *plausible* but rather in whether they are accurate in describing *actual* human behavior. There is now a growing body of evidence that documents the formal properties and efficiency of a number of fast-and-frugal heuristics (Czerlinski, Gigerenzer, & Goldstein, 1999; Goldstein et al., 2001; Martignon & Hoffrage, 1999) but this work needs to be complemented by empirical validation demonstrating that people do indeed use these heuristics in the environments in which they are claimed to operate. Without such empirical validation it seems that talk of a paradigm shift is premature.

In the two experiments reported here we examined the “take-the-best” (TTB) heuristic because of its status as one of the basic algorithms in the PMM framework (Gigerenzer & Goldstein, 1996). TTB is an extremely simple one-stop decision rule whose policy is “to take the best and ignore the rest” (Gigerenzer & Goldstein, 1996). In contrast to the more classically ‘rational’ decision strategies in which information is integrated (e.g., linear regression, Bayesian reasoning), TTB simply uses the ‘best’ piece of information applicable in a given situation. TTB operates according to two principles. The first—the recognition principle—states that in any given decision made under uncertainty, if only one amongst a range of alternatives is recognized, then the recognized alternative should be chosen (Goldstein & Gigerenzer, 2002). The second principle is invoked when more than one of the alternatives are recognized and the recognition principle cannot provide discriminatory information. In such cases, people are assumed to have access to a reference class of cues or features subjectively ranked according to their validities (where validity is defined as the probability that the feature will identify the correct alternative on a random selection of alternatives which differ on this feature). People are then thought to search the cues in descending order of feature validity until they discover a feature that discriminates one alternative from the others. Once this single discriminating feature has been found it is used to make a decision.

TTB is frugal because it bases its decision on a single feature, and fast because it does not bother searching through and integrating information about other alternatives (i.e., it is non-compensatory). Thus the TTB strategy can be thought of as comprising three basic building blocks: the search rule (search cues in order of descending validity), the stopping rule (stop after the first discriminating cue is discovered), and the decision

rule (choose the outcome pointed to by the first cue that discriminates).

The few studies that have sought empirical validation of TTB have found some evidence for its use (up to 65% of participants under conditions that favor the heuristic, e.g., Bröder, 2000, Experiment 4). In a previous set of experiments (Newell & Shanks, 2003) we used a share prediction task and created an experimental environment that we argued would promote adoption of the TTB strategy. The basic task required participants to predict which of two shares from two fictional companies was the more profitable. To help them make their predictions participants were able to buy information relating to the companies’ financial status. This information took the form of binary cues (e.g., Is it an established company?—YES/NO). On each trial participants were able to buy up to four pieces of information prior to making their prediction. This sequential buying of pieces of information—rather than the simultaneous and automatic provision of information used in some other tests of TTB (Bröder, 2000, Experiments 1 and 2; Jones et al., 2000)—allowed us to analyze participants’ information search and buying behavior. The aspect of participants’ behavior in which we were most interested was whether or not they continued to buy information *after* discovering a cue that discriminated between the two companies. We reasoned that seeking such ‘unnecessary’ information (i.e., information bought after the discovery of a discriminating cue) constituted a clear violation of TTB’s stopping rule. The stopping rule seems to us to be the most important of TTB’s components as it is the one which best distinguishes between the non-compensatory nature of the fast-and-frugal approach and the compensatory character of the more classic approaches to decision making such as linear regression or Bayesian calculation.

The experimental environments employed by Newell and Shanks (2003) were designed to promote the use of TTB. Participants were given a hint telling them the validities of the cues to encourage them to buy information in the order of the validities. The relative cost of information was high to discourage buying of unnecessary information. Furthermore, in one experiment a deterministic rather than a probabilistic environment was used to make it easier for participants to learn the relationships between the cues and the outcome. Despite these conditions contrived to promote the use of TTB, there was a high proportion of behavior that was inconsistent with the heuristic, particularly its stopping rule. In all the experiments a significant number of participants bought unnecessary information on some of the trials and in one experiment unnecessary information was bought on an average of 44% of the 120 test trials. We described this tendency to buy unnecessary information as a ‘weight of evidence’ strategy. Such

behavior is consistent with many other examples in the literature in which people are not always satisfied to make a decision on the basis of one piece of discriminating information but instead seek additional evidence to increase their confidence in their choice (Harvey & Bolger, 2001; Svenson, 1992, 1996).

It is worth noting that the experimental environments used by us and others to test fast-and-frugal heuristics (e.g., Bröder, 2000; Jones et al., 2000; Newell & Shanks, 2003) can be described as ‘menu-based’ inference tasks, because all the cue information required to make a decision is available (usually at some cost) on a ‘menu’ (i.e., on the computer screen or a piece of paper) in front of participants. These menu-based tasks contrast with ‘memory-based’ inference tasks in which cue information has to be retrieved from memory (Hastie & Park, 1986) have drawn a similar distinction between ‘on-line’ and ‘memory-based’ judgment tasks. Slegers, Brake, and Doherty (2000) have argued that fast-and-frugal heuristics are only applicable to memory-based inference tasks and therefore cannot be tested using menu-based tasks. Although the heuristics were designed for memory-based problems (such as the city population task, Gigerenzer & Goldstein, 1996) we argue that the explicit and precise descriptions of the heuristics lend themselves to being tested in experimental environments using a ‘menu’ of cues. Furthermore, proponents of the fast-and-frugal approach have used menu-based tasks to test its assumptions (e.g., Rieskamp & Hoffrage, 1999, see below).

Our results in conjunction with other empirical investigations of TTB (e.g., Bröder, 2000; Jones et al., 2000) demonstrate that TTB is clearly not universally adopted by participants—even under conditions strongly constrained to promote its use. We acknowledge that Gigerenzer and his colleagues do not assume that TTB is universal in the sense that it will be used by everyone regardless of the environment; in fact Gigerenzer is at pains to stress that the notion of an adaptive toolbox implies that the use of *different* heuristics is determined by the nature of the environment. Nevertheless, the failure to find any circumstances, however restricted, in which a particular heuristic (such as TTB) is almost universally employed seems to us to be problematic for the fast-and-frugal approach. It is problematic simply because if one cannot predict which heuristics will be used in which environments then determining the heuristic that will be selected from the toolbox for a particular environment becomes necessarily post hoc and thus the fast-and-frugal approach looks dangerously like becoming unfalsifiable. In this paper we ask whether an experimental environment can be constructed in which we might observe a higher proportion of participants adhering to the strategy, and thus provide further insight into the conditions under which particular heuristics are used.

One candidate task characteristic for promoting the use of TTB-consistent behavior was proposed by Rieskamp and Hoffrage (1999). They speculated that an increase in the complexity of a task would lead more participants to adopt simple non-compensatory strategies like TTB. Their speculation is based on the finding in a number of process-tracing studies that increasing the number of alternatives in a given decision domain and the number of dimensions or cues on which the alternatives are evaluated generally leads to the use of simplifying, non-compensatory strategies (e.g., Ford, Schmitt, Schechtman, Hults, & Doherty, 1989; Payne, 1976; Timmermans, 1993). In a review of 45 process-tracing studies, Ford et al. (1989) concluded, “Compensatory strategies were typically only used when the number of alternatives and dimensions were small” (p. 75). A clear prediction that follows from this conclusion is that increasing the complexity of the share task used by Newell and Shanks (2003) should provide an environment in which more people adopt non-compensatory strategies like TTB.

Our previous experiments and others that have employed similar tasks (e.g., Bröder, 2000; Jones et al., 2000) all used environments with four binary cues. In Experiment 1 we increased the number of cues to six. We reasoned that the addition of two cues should create an environment in which we were less likely to observe violations of the stopping rule of TTB. The key indicator of stopping rule violations was whether or not participants continued to buy information *after* discovering a discriminating cue. The experimental environment also allowed us to monitor participants’ adherence to TTB’s other two building blocks—the search rule (searching cues in order of descending validity) and the decision rule (choose the outcome pointed to by the first cue that discriminates).

2. Experiment 1

2.1. Participants

Twenty-four members of the University College London community took part in the experiment. Eleven were male and 13 were female. They had a mean age of 23.8 years (range 19–44, $SD = 5.5$). Most of the participants were undergraduate or graduate students at UCL.

2.2. Stimuli and design

The experiment was programmed in Microsoft Visual Basic 6 and run on an IBM compatible PC. Participants were presented with a series of two-alternative forced-choice decisions between the shares of two fictional companies (Share A and Share B). The shares were de-

scribed by six binary cues with semantic labels concerning aspects of the company's financial status. The cues had validities .90, .85, .80, .75, .70, and .65, where validity is defined as the probability that the cue will identify the correct alternative on a random selection of alternatives that differ on this cue. The assignment of the validities to the six nominal cues was counterbalanced across participants, but the screen position of the cues was constant. For each share there were 64 distinct cue patterns, and thus 2016 possible paired comparisons. A subset of 180 of these comparisons was randomly selected for each participant. For each comparison there was an associated probability of Share A being the more profitable. After each choice the probability that the chosen share was more profitable was computed according to Bayes' Rule, assuming conditional independence of the cues. It is true, of course, that conditional independence will not hold in some environments, but making the cues in our task conditionally independent facilitates testing the heuristic and interpreting the data, and for at least some real-world environments is probably a reasonable approximation. Appendix A provides details of the experimental environment and the calculations. A random number generator then determined which share was more profitable according to this probability.

2.3. Procedure

Participants were told that on each trial they would be asked to choose a share (one of two alternatives) that they thought would turn out to be the more profitable. To help them make this choice participants were able to buy up to six pieces of information about the two companies that the shares were from. The six pieces of information were: *Was the share trend positive over the last few months? Does the company have financial reserves? Does the company invest in new projects? Is it an established company? Is the company listed on the FTSE? and Is the employee turnover low?* Each piece of information cost 1p and was bought by clicking on a screen-button which revealed the value (YES or NO) of the piece of information for each share. Participants bought as much information as they desired and then selected either Share A or Share B. They were then told which share was most profitable for that trial and if they had made the correct prediction their account was incremented by 7p minus any money they had invested in buying information.¹

Participants were told that in making their predictions they should try to work out which pieces of information were most useful as not all the pieces were equally informative. To assist participants in learning

the relative usefulness of the cues a hint was provided after the first 60 trials and then again after 120 trials. This hint provided participants with the order of the usefulness of the six cues as determined by their actual validities (e.g., “‘Share Trend’ is the most useful piece of information followed by ‘Financial reserves,’ then ‘Invest in new projects,’” etc.). We note that ‘useful’ is a somewhat ambiguous term as it could apply either to the validity of a cue, the discrimination rate or a function of these two properties, namely ‘success’ (Martignon & Hoffrage, 1999).² However, the hint simply told participants that it would be more useful to them in terms of increasing their number of correct predictions to know the answer to the question, for example, “‘Is it an established company?’” than “‘Is it listed on the FTSE?’”. The hint was given twice at 60 trial intervals to encourage participants to follow the search rule by buying information in order of descending validity of the cues. There were 180 trials in total.

On completion of the 180 trials participants were asked to record their ratings (a number between 1 and 100) of the usefulness of each piece of information for helping them make their predictions. Finally, participants were given a post-test questionnaire asking them about the strategy they had used in making their predictions.

2.4. Results

2.4.1. Proportion correct

Fig. 1 shows the proportion of times the share chosen was the most profitable across the three 60 trial blocks. The overall mean proportion correct across the 180 trials was 70%, and above the chance level of 50% ($t(23) = 12.68, p < .001$). An analysis of variance revealed a significant linear trend across blocks of 60 trials ($F(1, 23) = 16.68, p < .001$).

2.4.2. On-line ranking and estimated usefulness of information

The computer recorded the order in which participants bought information. Each time information about the most valid cue was bought a ‘1’ was recorded, when information about the second most valid cue was bought a ‘2’ was recorded, and so on. This ‘on-line’ ranking of information buying behavior served as a manipulation check to establish whether participants bought information in the order of usefulness dictated by the hint, and thus followed TTB's search rule. The top row of Table 1 displays the mean on-line rankings derived from this measure collapsed across all 180 trials.

² Our recent investigations suggest that search in a similar task environment where no hint is provided is more closely associated with the order predicted by the success of the cues rather than their objective validity (Newell, Rakow, Weston, & Shanks, 2002).

¹ Note that UK1p = UK£.01 = US\$.02 = Eur .02 approximately.

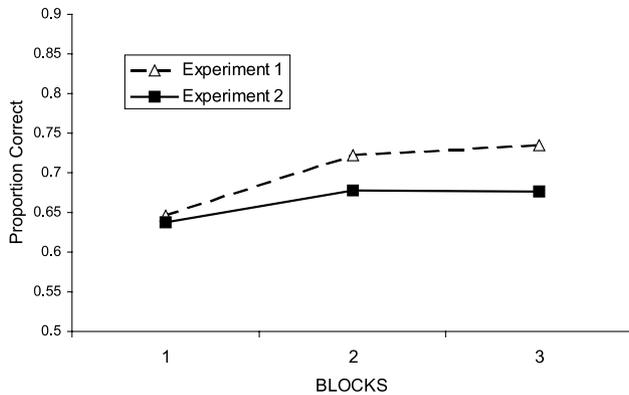


Fig. 1. Mean proportion of trials on which the 'correct' share was chosen in Experiments 1 and 2. *Note.* For Experiment 1 each block represents 60 trials, for Experiment 2 each block represents 40 trials.

To analyze this pattern we used a non-parametric Page Trend test for ordered alternatives (Siegel & Castellan, 1988). The Page Trend test evaluates the hypothesis that the means are ordered in a specific sequence versus the null hypothesis that the means are the same. The test indicated a significant trend ($z_L = 8.32$, $p < .0001$)³ across the six rankings in the predicted ascending order from the most valid to the least valid cue. A clearer picture of buying behavior was obtained by looking at rankings over the final 60 trials (after two hints had been provided) which are displayed in the second row of Table 1. As expected, the Page Trend test indicated a significant trend for the rankings on the final 60 trials ($z_L = 9.28$, $p < .0001$).

The estimated usefulness ratings of the six pieces of information were normalized by summing the ratings given to each cue and dividing each rating by this total for each participant. The mean, normalized, usefulness ratings for the six pieces of information are displayed in the third row of Table 1. A Page Trend test for ordered alternatives indicated a significant trend ($z_L = 8.51$, $p < .0001$) across the six ratings in the predicted descending order from the objectively most useful to the objectively least useful piece of information.

2.4.3. Analysis of TTB's three components

Previous research indicates that large individual differences in strategy use exist for these types of task (Bröder, 2000; Newell & Shanks, 2003) and thus we examined group and individual data patterns for the three components of TTB—search, stopping, and decision. Participants needed to learn about the structure of the environment prior to receiving any hints about the validities of cues. This learning may have led to fluctuations in behavior that are not accurate representations

of participants' strategies. For this reason we analyzed behavior during only the final 60 trials. In this final stage after 120 trials and two hints we reasoned that any violations of the heuristic's components were less likely to be attributable to exploration of the experimental environment. Table 2 displays individual data for Experiment 1 for these final 60 trials.

2.4.4. Search rule

The on-line ranking and estimated usefulness measures illustrate that at a group level participants behaved in a way that is consistent with the underlying assumptions of the search rule of TTB—namely that they had access to a hierarchy of cues that they subjectively ordered correctly and that they searched through in order of descending validity. The second column of Table 2 indicates that this search behavior was fairly consistent across participants. Recall that the computer recorded a "1" every time the most valid piece of information was bought, a "2" for the second most valid and so on. Thus the values in the second column of Table 2 are the mean rankings for the piece of information bought first, second, third, fourth, etc., across the final 60 training trials. Only two participants (p4 and p10) failed to buy the most valid piece of information first, and 18 participants (75%) correctly bought the most valid first followed by the second most valid piece of information.

2.4.5. Stopping rule

For the stopping rule we wanted to test was whether search was terminated after a participant found a cue that discriminated between the alternatives. In order to test whether participants adopted this strategy the computer recorded the number of trials on which information was bought after a discriminating cue was identified. This allowed us to identify trials on which a participant bought a discriminating cue (e.g., 'YES' for Share A and 'NO' for Share B) and then went on to buy 'unnecessary' information about another cue.

The mean proportion of unnecessary information bought was .20 ($SE = .05$). Thus at a group level, participants' behavior was inconsistent with the stopping rule, because strict adherence to the strategy would predict a value of 0 for the unnecessary information measure. However, the third column of Table 2 indicates that in fact eight participants (p7, p10, p12, p13, p15, p18, p23, and p24) out of 24 (33%) did not buy any unnecessary information. These participants thus completely adhered to the stopping rule. In contrast, five participants bought unnecessary information on over 50% of trials.

2.4.6. Decision rule

For testing the decision rule we were interested in the proportion of trials on which unnecessary information was bought, and participants' decisions were *not* in the direction pointed to by the first discriminating cue that

³ We report the z_L value for the Page Trend tests following the conversion formula for an N exceeding 12 (Siegel & Castellan, 1988).

Table 1
Experiment 1

Cue validity	Most		Least			
On-line ranking (<i>SD</i>)—all trials	2.08 (1.0)	2.67 (.49)	3.16 (.49)	3.98 (.71)	4.33 (.66)	4.49 (1.1)
On-line ranking (<i>SD</i>)—last 60 trials	1.29 (.68)	2.18 (.51)	2.94 (.40)	4.32 (.81)	4.78 (.66)	5.16 (1.1)
Estimated usefulness (<i>SD</i>)	.30 (.15)	.22 (.07)	.18 (.08)	.12 (.08)	.09 (.06)	.09 (.08)

Top row: on-line ranking of the six cues over all 180 trials. Second row: on-line ranking of the six cues over the last 60 trials. Third row: Normalized Estimated Usefulness ratings of the six cues.

Table 2
Individual participant data for Experiment 1 (last 60 trials)

Participant No.	On-line ranking	Unnecessary information	Decision rule violations	No information (guessing)
1	1-2-3-6-5-4	.23	.07	0
2	1-2-3-5-●-●	.65	.31	0
3	1-2-3-4-5-6	.03	0	0
4	2-3-4-1-5-6	.75	.27	0
5	1-2-3-5-4-●	.71	.06	.25
6	1-3-2-6-4-5	0	0	0
7	1-2-3-●-●-●	0	0	0
8	1-2-3-4-5-6	.02	0	0
9	1-2-3-4-5-6	.56	.12	0
10	3-1-2-5-4-6	0	0	0
11	1-2-3-●-●-●	.33	.50	.90
12	1-2-3-4-5-6	0	0	.02
13	1-2-●-●-●-●	0	0	0
14	1-2-3-●-●-●	.10	.17	0
15	1-2-3-4-5-6	0	0	0
16	1-2-3-4-5-6	.36	.30	.07
17	1-2-3-4-5-6	.52	.13	0
18	1-2-3-●-●-●	0	0	.65
19	1-2-6-3-4-5	.02	0	0
20	1-3-2-4-●-●	.12	.29	0
21	1-3-5-6-2-4	.16	.29	.28
22	1-2-6-3-4-5	.30	.06	0
23	●-●-●-●-●-●	0	0	1.0
24	1-2-3-4-●-●	0	0	0

Note. On-line ranking is derived from the order in which participants bought information and thus provides an indication of whether participants followed the search rule. A ● indicates that no further pieces of information were bought. Unnecessary Information refers to the proportion of trials on which information was bought after discovering a discriminating cue (violating the stopping rule). This measure is conditionalized on whether any information was bought at all. For example p5 guessed on 25% of trials (15 out of 60); on the remaining 75% of trials (45 out of 60) she bought unnecessary information on 71% of trials (32 out of 45). ‘Decision rule violations’ refer to the proportion of trials on which unnecessary information was bought and participants’ choices were *not* in the direction pointed to by the first discriminating cue that they purchased. No information or ‘guessing’ refers to the proportion of trials on which participants did not purchase any information prior to making their decision.

they purchased. For example, if a participant bought the most valid cue and it discriminated (e.g., ‘YES’ for Share A and ‘NO’ for Share B) but then went on to buy further cues, and finally chose Share B this constituted a violation of the decision rule (and the stopping rule). Note that this measure does not distinguish between trials on which the cues acquired after the first discriminating cue point in the opposite direction, the same direction, or are tied (i.e., both ‘YES’ or both ‘NO’).

The mean proportion of trials on which there were decision rule violations was .11. Thus at a group level, participants’ behavior was somewhat inconsistent with the decision rule, because strict adherence to the strategy would predict a value of 0 for this measure. However, at an individual level, the fourth column of Table 2 shows that 12 (p3, p6, p7, p8, p10, p12, p13, p15, p18, p19, p23,

and p24) out of 24 participants (50%) never violated the decision rule.

2.4.7. Classification of participants’ strategies

We have argued that the measure which distinguishes best between fast-and-frugal behavior and a more compensatory mechanism, is the acquisition of ‘unnecessary’ information. For this reason, we decided to classify participants’ strategies on the basis of this measure. This resulted in the identification of three separate strategies: a “frugal” strategy in which a decision was made after discovering one discriminating cue, a “weight of evidence” strategy in which participants continued to buy ‘unnecessary’ information after discovering a discriminating cue, and “guessing” in which no information was bought at all.

2.4.8. Frugal strategy

Table 2 indicates that 11 (p3, p6, p7, p8, p10, p12, p13, p14, p15, p19, and p24) out of 24 participants (46%; 95% $CI = \pm 20\%$) appear to have adopted a frugal stopping rule. Some of these participants did still buy some unnecessary information, but the small amounts (a maximum of 10% by p14) could arguably be attributed to response error or attention lapses.⁴ Participants in this behavioral category bought an average of 1.77 pieces of information per trial (out of a maximum of 4) and earned an average of £5.92, in addition to their participation fee, and their mean proportion correct across all 180 trials was 73%.

2.4.9. Weight of evidence

Table 2 shows that nine (p1, p2, p4, p5, p9, p16, p17, p20, and p22), out of the 24 participants (38%; 95% $CI = \pm 19\%$) adopted what we have termed a weight of evidence strategy. On a significant proportion of trials (mean 44%, minimum 12%, and maximum 75%) these participants bought more information after discovering a discriminating cue. For these participants, presumably basing their decisions on what for certain trials was only one sixth of the available information was perceived as too risky. However, it is not the case that these participants bought all available information on each trial. The mean amount of information bought per trial was 2.55, and the mean amount of unnecessary information was 1.26. The average amount earned by participants in this category was £4.36 and their mean proportion correct across all 180 trials was 71%.

2.4.10. Guessing

Four participants, (p11, p18, p21, and p23), (16%; 95% $CI = \pm 15\%$) guessed on a considerable proportion of trials. Three other participants guessed on some trials, notably p5, but in each case it was not their dominant strategy. These participants earned an average of £6.07 and their mean proportion correct across all 180 trials was 58%.

2.4.11. Post-test questionnaires

In addition to behavioral data we collected participants' reports of the strategy they thought they had used. In general, self-reports matched behavior, but a number of inconsistencies were apparent. Therefore, we

grouped participants solely on their behavior as this provided a less ambiguous classification of the actual strategy adopted.

2.5. Discussion

We examined the possibility that increasing the number of cues in the share prediction task would create an experimental environment in which violations of the TTB strategy were minimized. The results of Experiment 1 clearly demonstrate that using an environment with six cues did not eliminate violations of TTB. If adherence to each rule (i.e., validity ordered search for the first two cues, and values for the stopping, decision and guessing measures within the 10% margin for attention lapses, see footnote 4) is demanded then eight participants (p3, p7, p8, p12, p13, p15, p19, and p24) (33%) can be classified as displaying behavior consistent with TTB.

Although the TTB strategy may not have fared too well in our environment, the fast-and-frugal approach in general did well. The majority of participants displayed validity ordered search, at least for the first two or three cues. Furthermore, the largest proportion of participants were those adopting a frugal stopping rule, who in addition to stopping search after acquiring only one piece of discriminating information, conformed to the decision rule by always choosing the alternative pointed to by the discriminating cue (with the exception of p14 who chose in the opposite direction on 17% of trials).

However, almost 40% of participants tended to buy more information than was necessary—even when it was cognitively (in terms of effort) and financially disadvantageous to do so, thus violating the stopping rule. These participants all violated the decision rule too, on at least some of the trials (minimum .06, maximum .30, and mean .18). We classified these participants as adopting a weight of evidence strategy. This strategy appears not to be compensatory in the sense that a classical weighted additive rule is because participants did not buy *all* the information available on each trial. It seems rather that these participants wanted to accumulate one extra piece of information per trial on average, *after* discovering a discriminating cue, perhaps in order to reduce their uncertainty about the outcome of, or the risk associated with, their choice.

Finally, a small proportion of participants chose to guess on a large proportion of trials. Although this guessing strategy was perhaps economically adaptive in this environment (this group had marginally higher earnings than the frugal strategists) it was not a successful strategy for learning about the usefulness of the cues or increasing the number of correct predictions made. Note that both these objectives were explicitly stated in the instructions, and this is probably why only a minority of participants chose to adopt a guessing strategy. It is also of course possible that for these few

⁴ It is unrealistic to assume that the rule is completely error-free and therefore unnecessary information will sometimes be bought by mistake. To estimate the frequency of such occasions we calculated the proportion of trials on which unnecessary information was bought by those participants in the two experiments whose individual on-line ranking, estimation and self-report data were completely consistent with the systematic use of a frugal stopping rule. These eight participants bought unnecessary information on an average of 1% of trials ($SD = 2\%$). Thus the upper limit of 10% that we used as a 'cut-off' point is a generous estimate.

participants the pay-off structure of the environment was not sufficient to engage their critical thinking.

The pattern of results leaves open the question whether any environment would favor the TTB strategy to the extent that all, or almost all of the participants would conform to each of its components. In Experiment 2 we approached this question by designing what we considered to be the simplest possible situation in which TTB might operate.

3. Experiment 2

TTB is a decision heuristic used to make choices between two alternatives. When information about the two alternatives is known, a decision is made on the basis of the most valid piece of information that discriminates between the alternatives—without seeking any further information. Thus, the simplest situation in which TTB can be tested is one in which there are two alternatives and two cues that differ in validity. TTB predicts that the less valid cue should *only* be examined when the most valid cue does not discriminate. This prediction would be especially strong in a situation in which participants know the validities of the cues and where the relative cost of information is high. Experiment 2 used such an environment. We reasoned that the simplicity and the transparency of the environment in terms of the relationship between the cues and the outcome would reduce the temptation for participants to seek further evidence after discovering a discriminating cue.

Of course, one could argue on the basis of the evidence from the process-tracing literature that the reduction to two cues should favor non-frugal strategies and that TTB use will be low in such environments. However, Experiment 1 illustrated that behavior in the share prediction task was not congruent with the findings in the process-tracing literature—we found that a significant proportion of participants used non-frugal strategies when the environment was relatively complex. Given this incongruence, we felt confident that the two cue environment would promote rather than discourage TTB use. Furthermore, we speculated that some participants in Experiment 1 might have been reluctant to make their decisions on the basis of one piece of discriminatory information because of the potential impact on the outcome of up to five remaining pieces of information. With only one extra piece of information to buy, would participants still be tempted to seek further evidence after discovering a discriminating cue?

3.1. Participants

Twenty-four members of the University College London community took part in the experiment. Nineteen were female and five were male. They had a mean

age of 22.5 years (range 18–38, $SD = 4.7$). The majority were undergraduate or graduate students at UCL.

3.2. Stimuli, design, and procedure

Experiment 2 used the same share prediction task as Experiment 1 with some modifications. Information concerning aspects of the company's financial status was described by only two, rather than six, binary cues, which had validities of .80 and .60. For each share there were 4 distinct cue patterns and thus 12 possible paired comparisons. Each comparison was shown ten times in a different random order for each participant for a total of 120 trials. The two pieces of information were: *Does the company invest in new projects?* and *Is it an established company?* A hint telling participants the order of usefulness of the two pieces of information was provided after 30 trials. If the correct share was chosen the private account was incremented by 5p minus any money that had been invested to buy information. As in Experiment 1 each piece of information cost 1p. Thus there was a slight increase in the relative cost of information from a seventh of the potential profit in Experiment 1 to a fifth in this experiment.

3.3. Results

3.3.1. Proportion correct

Fig. 1 shows the proportion of times the share chosen was the more profitable one over the 120 trials. The mean was 66%. This proportion was well above the chance level of 50% ($t(23) = 11.17, p < .001$). However, an analysis of variance did not reveal a significant linear trend across the three blocks ($F(1, 23) = 2.31, p = .14$).

3.3.2. On-line ranking and estimated usefulness of information

Table 3 displays the mean on-line rankings over the total 120 trials and the last 30 trials and the estimated usefulness of each piece of information. A Wilcoxon test revealed a significant difference between the two rankings ($z = 2.86, p < .01$) and between the usefulness ratings of the two cues, ($z = 3.67, p < .001$). The pattern for these measures indicates that participants both subjectively rated and searched through cues in the correct order. This behavior is even more apparent for the rankings of

Table 3
Experiment 2

Cue validity	Most	Least
On-line ranking (SD)—all trials	1.19 (.25)	1.76 (.33)
On-line ranking (SD)—last 30 trials	1.13 (.31)	1.86 (.33)
Estimated usefulness (SD)	.71 (.19)	.29 (.19)

Top row: On-line ranking of the six cues over all 120 trials. Second row: On-line ranking of the six cues over the last 30 trials. Third row: normalized estimated usefulness ratings of the two cues.

the final 30 trials (second row of Table 3), and as expected the Wilcoxon test revealed a significant difference between the rankings ($z = 3.28, p < .002$).

3.3.3. *Analysis of TTB's three components*

We examined trials from the latter stages of the experiment, concentrating on the last block of 30 trials after participants had received a hint and had had 90 trials to learn about the experimental environment.

3.3.4. *Search rule*

The second column of Table 4 displays individual participant data for on-line ranking—our indicator of search rule adherence. As the group analysis suggested, all but two participants (p7 and p10) followed the search rule by buying the most valid piece of information first.

3.3.5. *Stopping rule*

The mean proportion of unnecessary information bought was .11. Thus again at a group level performance did not strictly adhere to the stopping rule. However, the third column of Table 4 indicates that an impressive 15 participants (p1, p2, p3, p4, p7, p8, p12, p13, p14, p15, p16, p20, p21, p22, and p23) out of 24 (62%) never violated the stopping rule.

3.3.6. *Decision rule*

The mean proportion of trials on which the decision rule was violated was .11. Thus again at a group level performance did not strictly adhere to the rule. How-

ever, the fourth column of Table 4 shows that a very high number of participants, 17 (p1, p2, p3, p4, p5, p7, p8, p12, p13, p14, p15, p16, p19, p20, p21, p22, and p23) out of 24 (71%) never violated the decision rule.

3.3.7. *Classification of participants' strategies*

Consistent with the analysis of Experiment 1 we focused on participants' patterns of information acquisition as a key indicator of the type of strategy being used. Participants were classified as adopting a frugal, a weight of evidence, or a guessing strategy with regard to information acquisition.

3.3.8. *Frugal strategy*

As the stopping rule analysis above indicated, 62% (95% *CI* = ±19%) of participants appear to have used a frugal stopping rule. Consistent with Experiment 1 some participants (p5 and p6) did buy some unnecessary information on some trials (7%) but at levels arguably low enough to be attributed to response error or attention fluctuations (see footnote 4). Participants who adopted a frugal strategy bought an average of 1.15 pieces of information per trial, earned on average £2.59 in addition to their participation fee and their mean proportion correct across all 120 trials was 67%.

3.3.9. *Weight of evidence*

Table 4 shows that six (p9, p10, p11, p17, p18, and p24) out of 24 participants (25%; 95% *CI* = ±17%) adopted a weight of evidence strategy seeking further

Table 4
Individual participant data for Experiment 2 (last 30 trials)

Participant No.	On-line ranking	Unnecessary information	Decision rule violations	No information (guessing)
1	1-2	0	0	.50
2	1-●	0	0	0
3	1-2	0	0	0
4	1-●	0	0	0
5	1-2	.07	0	0
6	1-2	.07	.50	0
7	2-●	0	0	0
8	1-2	0	0	0
9	1-2	.20	.67	0
10	2-1	.70	.48	0
11	1-2	.67	.05	0
12	1-2	0	0	0
13	1-●	0	0	.53
14	1-2	0	0	.03
15	1-2	0	0	0
16	1-●	0	0	0
17	1-2	.26	.25	0
18	1-2	.11	.67	.10
19	1-2	.10	0	0
20	1-●	0	0	.01
21	1-●	0	0	.23
22	1-2	0	0	0
23	1-●	0	0	0
24	1-2	.48	.14	.03

Note. See note accompanying Table 2 for a description of each measure.

information after discovering a discriminating cue. Participants 10 and 11 did in fact, buy the unnecessary information on every trial in which it was possible to do so. Note that out of the 12 possible comparisons in our experimental environment, on four there was no unnecessary information to buy as the most valid cue did not discriminate, thus the maximum value for our measure across all 120 trials was .66. In the final 30 trials the mean amount of unnecessary information bought by these participants was 40% (maximum 70%, minimum 10%). (Note that the maximum was 70 and not 66% due to sampling from the final 30 trials). Participants who adopted the weight of evidence strategy bought 1.62 pieces of information per trial, and 1.0 pieces of unnecessary information, on average. They earned an average of £2.19 and their mean proportion correct across the 120 trials was 68%.

3.3.10. *Guessing*

Three participants (p1, p13, and p21), (13%; 95% $CI = \pm 13\%$) chose not to buy any information on a significant proportion of trials. These participants earned an average of £2.59 and their mean proportion correct across the 120 trials was 58%.

From this simplified, two-cue version of the task, we were able to identify a further group of participants who guessed in a slightly different manner. In addition to guessing on the basis of no information, participants could guess after buying a non-discriminatory piece of information, rather than going on to buy another piece of information as TTB predicts. Out of the 12 possible paired comparisons, there were four on which the first cue did not discriminate. An examination of individual data over the last 30 trials revealed that seven participants (p2, p3, p4, p12, p20, p21, and p23) simply guessed whenever the more valid piece of information did not discriminate. They did this on 100% of trials that it was possible to do so (in the last 30 trials, there were on average 9 trials where the first cue did not discriminate). Therefore, these participants were not strictly adhering to the particular type of ‘frugal’ strategy to which they had previously been classified (see above). They made a decision to only buy one piece of information on every trial, and made their choice regardless of whether this piece of information was discriminatory.

Due to the large number of cues in Experiment 1, and hence a more complex environment, this behavior could not be clearly identified. When presented with six cues it would be possible for participants to both buy unnecessary information and make their choice after buying a non-discriminatory cue on the same trial, thus making it very difficult to classify behavior.

3.3.11. *Post-test questionnaires*

Consistent with Experiment 1, verbal report data indicated some discrepancies between behavior and re-

ported strategies. For this reason we again based strategy classifications purely on the behavioral data.

3.4. *Discussion*

We suggested that a simple and transparent environment in which there were only two cues would favor a fast-and-frugal strategy like TTB. Eight participants (p3, p5, p8, p12, p14, p15, p19, and p22) (33%) behaved in a manner consistent with all of TTB’s building blocks (validity ordered search, and violations of the stopping, decision and guessing measures within a 10% margin). This is the same proportion as observed in Experiment 1. Consistent with Experiment 1, the remaining two-thirds of participants displayed behavior that violated at least one of TTB’s rules, including a sizable proportion that used a weight of evidence strategy.

It was the case, however, that the majority of participants used a frugal stopping rule. This frugal rule either ended search after discovering a discriminating cue, as predicted by the TTB heuristic, or terminated search after just looking at the most valid cue—regardless of whether it discriminated. This second strategy could be described as the ‘‘Take-One’’ heuristic (Gigerenzer, personal communication), in which the stopping rule states ‘‘stop after the most valid cue is bought—regardless of whether it discriminates.’’ We acknowledge that guessing—either on the basis of no information or non-discriminatory information—is both fast and frugal and therefore consistent with the framework.

4. **General discussion**

The aim of this research was to create an environment in which violations of the TTB strategy that we had observed previously (Newell & Shanks, 2003) were eliminated or reduced. We attempted to achieve this aim by manipulating the number of cues or pieces of information in the experimental environment. In Experiment 1 we increased the number of cues to six (from four that we and others had used previously) in the hope that increasing the complexity of the task would promote the adoption of non-compensatory strategies like TTB (cf. Ford et al., 1989; Payne, 1976; Rieskamp & Hoffrage, 1999). Despite using a six-cue environment, only eight participants out of 24 were classified as adhering to each of TTB’s rules (search, stopping, and decision). Furthermore, although a fairly large proportion of participants adopted frugal strategies (46%) a similar number (38%) adopted a non-frugal weight-of-evidence strategy.

In Experiment 2 we tried to design the simplest situation in which TTB can operate—a two alternative, two-cue environment. We reasoned that the simplicity and transparency of the environment would promote TTB use to the point that violations might be elimi-

nated. In this environment, again only eight out of 24 participants behaved in a manner completely consistent with each of TTB's rules, but overall the adoption of frugal strategies clearly outnumbered the use of non-frugal strategies (62 and 25% of participants, respectively).

We noted in Section 1 that the fast-and-frugal approach has become increasingly popular in recent years and has been applied to a variety of domains (e.g., Dhimi & Ayton, 2001; Elwyn et al., 2001; Seeley, 2001). It is important to emphasize that our results do not challenge this popularity in terms of the applicability of the approach. The results do not undermine the power of the TTB heuristic. Indeed those participants who did adopt the heuristic or components of it earned more money than those who used a less frugal strategy. Our results simply help us to understand the type of conditions under which we might expect people to adopt particular heuristics—this we believe is the most important challenge facing the fast-and-frugal approach. There is a strong need to complement research showing how effective heuristics can be with research showing that they provide an accurate description of *actual* human behavior.

4.1. Violations of TTB's 'building blocks'

A conservative estimate of TTB-consistent behavior across both experiments suggests that 16 participants out of 48 (33%) conformed to all three of TTB's explicitly defined building blocks, even allowing for a degree of lapses in attention (see footnote 4). The presence of a high proportion of deviations from TTB (be they frugal or non-frugal) suggests that either the heuristic is only adopted by a minority of the population and thus that its psychological reality is doubtful, or that we are still just 'looking in the wrong place' for the conditions that promote its use. Obviously no test of a theory can be exhaustive and we stress that we are only testing the assumptions of the fast-and-frugal approach for this particular type of decision task, but at the very least our results suggest that the description of TTB as it stands is too restrictive to encompass the wide variability in the behavior we observed.

For example, the description of the heuristic only permits guessing when there are either no cues available in the environment or when all the cues present are non-discriminatory. In contrast to these situations we observed participants who chose to guess rather than purchase remaining discriminating cue information (the Take-One heuristic). It is relatively easy to think of other situations in which the current description of TTB is rather restrictive. Consider an experimental environment identical to the one employed here, but in which making the correct choice on the final trial determines whether or not the participant receives all the money

accumulated throughout the experiment. It seems likely (although we acknowledge that it is an empirical speculation) that the participant would want to be as sure as possible about being correct and would therefore buy all the information available and not just one discriminating piece.

This issue of the severity of the outcome of decisions is one that needs to be addressed within the fast-and-frugal framework. It may be that when an outcome is severe a different heuristic or a combination of heuristics are selected from the 'toolbox' but at this stage it is not clear how the environment and/or a homunculus predicts or controls this selection process—a point that has not passed unnoticed (Goldstein et al., 2001). As we emphasized in the introduction, unless one can a priori specify the conditions under which certain heuristics will be selected over others—and *provide human data consistent with these predictions*—the predictive and explanatory power of the fast-and-frugal approach remains questionable.

4.2. Ecological analysis of environments

One method for predicting when particular heuristics might be used is to analyze a task environment from an ecological perspective. One potential criticism of the experiments reported here is that we relied on the findings from other process-tracing studies to inform our predictions about whether or not we would observe TTB-consistent behavior. Although this seems to us to be a legitimate way to conduct research, another approach is to create an environment and then to determine the ecologically rational heuristics for that environment.⁵ For the sake of simplicity we will consider the environment used in Experiment 2, though the same rationale holds for Experiment 1.

The two cues in Experiment 2 had validities of .80 and .60; one cue had a discrimination rate of .66 and the other a discrimination rate of .33. Each cue cost 1p, which was deducted from the maximum winnings of 5p. Therefore the maximum expected winnings for a TTB strategy can be calculated as follows:

$$\begin{aligned} E(\text{TTB}) &= .80 \times .66 \times (5 - 1)p + .60 \times .33 \times (5 - 2)p \\ &= \text{£}2.73p. \end{aligned}$$

A similar analysis for the Take-One (TO) heuristic demonstrates that the expected winnings through using this heuristic is the sum of two terms, the expected winnings from buying the most valid cue and the expected winnings from guessing:

$$\begin{aligned} E(\text{TO}) &= .80 \times .66 \times (5 - 1)p + .50 \times .33 \times (5 - 1)p \\ &= \text{£}2.80p. \end{aligned}$$

⁵ We thank Gerd Gigerenzer for suggesting this analysis.

Finally, a simple guessing strategy in which no information is bought at all has expected winnings of:

$$E(\text{guess}) = .50 \times 5p = \text{£}2.50p.$$

This kind of ecological analysis suggests that TTB is not the only heuristic that we should observe in the environment used in Experiment 2. Indeed the closeness of the expected winnings for the three strategies suggests that we should observe all the heuristics with guessing being the least frequent—which is the pattern we found.

However, what such an ecological analysis does not predict—and can never predict in the non-compensatory environments used in the current experiments and previously (Newell & Shanks, 2003)—is the existence of a weight of evidence strategy of the kind adopted by a significant proportion of participants. An ecological analysis is undoubtedly a useful tool in helping to establish when particular heuristics might be used, but the presence of the weight of evidence strategists highlights that such an analysis is unable to predict the full variability present in actual human behavior.

It is worth clarifying here that by using the label ‘weight of evidence’ we are *not* implying an analytic strategy in which each piece of information is obtained and given a relative weight (e.g., a type of weighted additive rule). The strategy we observed was not one in which a participant bought every available piece of information, it was simply one that involved obtaining, on average, one extra ‘unnecessary’ piece of information before making a decision. Note also that obtaining this unnecessary information did not necessarily imply that the final decision rule of TTB was violated. As indicated in Tables 2 and 4, many participants in both experiments bought information after discovering a discriminating cue but ultimately chose the option pointed to by the *first* discriminating cue that they had discovered (as predicted by the decision rule). Although behavior of this kind is not a violation of the decision rule, accumulating *more* than one piece of discriminating information before making a decision does not seem, to our minds at least, to be synonymous with *one-reason* decision making.

We speculate that the extra piece of information served to reduce participants’ uncertainty about an outcome or the risk associated with a choice. Evidently this element of risk assessment or uncertainty reduction is an important factor for a significant proportion of participants. Another possible reason for the accumulation of extra information is to increase the amount of differentiation between the two alternatives. Svenson (1992, 1996) cites evidence in support of a ‘differentiation and consolidation’ theory, which proposes that people continue to acquire information until a particular candidate is “sufficiently superior for a decision” (Svenson, 1992, p. 143). Montgomery and colleagues (e.g., Montgomery, 1983; Montgomery & Willén, 1999)

propose a similar idea in the ‘search for dominance structure’ theory in which information is structured and restructured “in such a way that one alternative becomes the self-evident choice” (Montgomery & Willén, 1999, p. 148). Although these theories are typically applied to situations involving a greater number of alternatives and attributes (e.g., Harvey & Bolger, 2001) than in the experiments presented here, it is clear that they support the conjecture that people are not always satisfied to make a decision on the basis of one piece of discriminating information.

4.3. *Half full versus half empty or a fundamental problem?*

Throughout this article and previously (Newell & Shanks, 2003), we have tended to highlight participants’ violations of the component rules of particular heuristics rather than adherence to those rules. This focus lays us open to the criticism that we are being too negative in our assessment of the fast-and-frugal approach. It could be said that we are taking “the glass is half empty” view rather than emphasizing that it is “half full.” However, we argue that our position does not reflect a simplistic negative bias, but rather that the violations of the heuristics are more ‘newsworthy’ because they represent a fundamental problem for the fast-and-frugal approach.

In Experiment 1, at a group level, participants adhered to the search, stopping and decision rules in 75, 80, and 89% of all possible cases respectively. For Experiment 2 the figures are even more impressive: 92% for the search rule and 89% for the stopping and decision rules. Despite the fact that these group means belie wide individual variability (see Tables 2 and 4) these results appear extremely good for the fast-and-frugal approach—especially as they were achieved in an environment with no adjustable parameters or error terms.

However, this latter point is perhaps both the most appealing and yet potentially the most problematic aspect of the fast-and-frugal approach. The determinism of the rules is appealing because it makes them simple and thus supposedly more psychologically plausible than more computationally complex decision models (though see Chater, Oaksford, Nakisa, & Redington, in press). But because the heuristics are described deterministically there is no room for error and therefore any deviations from the predicted behavior are necessarily problematic and can only be explained by the introduction of another heuristic or building block (and another, and another—potentially). At this point the division between the simple, transparent heuristic, and a complex, opaque model—a division much championed by Gigerenzer and colleagues—becomes rather muddled. If the ‘free parameters’ are the invention of additional heuristics, why is this preferable to a model that incorporates free parameters and by doing so achieves wider predictive ability—e.g., PROBEX (Juslin & Persson, 2002)?

At the individual level we have found repeated violations of the behavior predicted by simple heuristics—including violations that are not predicted by an ecological analysis of the environment. As we suggested in Section 1, if the only way to explain these violations within the fast-and-frugal framework is through the post hoc invention of new heuristics or building blocks, then the framework begins to appear dangerously unfalsifiable.

4.4. Task characteristics or personality characteristics?

Our focus in these experiments has been on the characteristics of the task, but it is quite likely that personality characteristics will have an effect on the adoption of particular behavioral strategies. One aspect of the data that supports this conjecture is the fact that although we have varied a range of task characteristics across the experiments described here and previously (Newell & Shanks, 2003), the relative proportions of different strategy use is fairly constant. At present our results do not allow us to predict what types of people are likely to be ‘fast-and-frugal’ and what types are likely to be more conscientious evidence gatherers. Future investigations could combine experimental tasks of the kind reported here with assessments of personality or decision making styles or a simple measure of risk aversion in order to gain a clearer picture of not only when and why fast-and-frugal heuristics might operate but also the type of people who we might expect to make use of them. For example, we might expect a participant who is classified as a ‘weight of evidence’ strategist to load more heavily on the conscientiousness component of the Five Factor model of personality than a TTB user (e.g., Wiggins, 1996). Another potential tool is Nygren’s Decision-Making Styles Inventory, which classifies participants as having either analytic or intuitive decision making styles (e.g., Nygren & White, 2001). Presumably TTB users would be classified as ‘intuitive’—basing their decisions on less information than the more ‘analytic’ weight of evidence strategists. Similarly, a TTB user might be less risk averse in a typical gamble problem (e.g., Kahneman & Tversky, 1979) than a weight of evidence strategist. Recent interesting work by Bröder (2002) has begun to examine these issues, and although he failed to find any personality measures that correlated with strategy use, he found evidence that people who used TTB in non-compensatory environments were more intelligent than those who employed compensatory strategies.

5. Conclusion

In his paper ‘Rational choice and the structure of the environment,’ Simon (1956) concluded that his approach to the description of rational behavior was in

“closer agreement with the facts of behavior as observed in laboratory and field” (p. 138) than that of more traditional economic and statistical theories. Before the current extensions of Simon’s ideas as expounded by Gigerenzer and colleagues (e.g., Gigerenzer & Todd, 1999) induce any paradigm shifts in judgment and decision-making, empirical validation of the fast-and-frugal heuristics is highly necessary. Our results suggest that the approach is yet to achieve ‘close agreement’ with behavior observed in the laboratory.

Acknowledgments

The support of the Economic and Social Research Council (ESRC) and The Leverhulme Trust is gratefully acknowledged. The work was part of the programme of the ESRC Research Centre for Economic Learning and Social Evolution (ELSE), University College London. We thank Gerd Gigerenzer, Mike Doherty, and one anonymous reviewer for their helpful and insightful comments on an earlier version of this manuscript.

Appendix A

Probabilities that Share A or B is more profitable for the different cue patterns, and the likelihood ratio (LR) favoring A over B in Experiment 1

Cue	Pattern	$P(A = 1/C_i)$	$P(B = 1/C_i)$	LR (A:B)
C ₁	10	.90	.10	9/1
	01	.10	.90	1/9
C ₂	10	.85	.15	85/15
	01	.15	.85	15/85
C ₃	10	.80	.20	4/1
	01	.20	.80	1/4
C ₄	10	.75	.25	3/1
	01	.25	.75	1/3
C ₅	10	.70	.30	7/3
	01	.30	.70	3/7
C ₆	10	.65	.35	65/35
	01	.35	.65	35/65

Note. C₁–C₆ refer to the individual cues (e.g., Established company?). A cue is diagnostic only when there is a YES for one share and a NO for the other share. The Pattern column displays the two diagnostic patterns for each of the four cues (1 = YES, 0 = NO). $P(A = 1/C_i)$ refers to the probability that Share A is more profitable (i.e., A = 1) for a given cue pattern; $P(B = 1/C_i)$ refers to the probability that Share B is more profitable (i.e., B = 1) for a given cue pattern.

To calculate the likelihood of a share from company A being more profitable than one from company B, the likelihood ratios corresponding to the particular

patterns are multiplied. For example, if the pattern for company A was 11111 and for company B was 00000, then all six cues would be diagnostic and the share from company A would be $(9/1 \times 85/15 \times 4/1 \times 3/1 \times 7/3 \times 65/35) = 2638$ times more likely to be more profitable. In contrast, if the patterns were 000001 for company A and 000000 for company B only C_6 would be diagnostic and the share from company A would be 65/35 or approximately twice as likely to be the more profitable. The method for calculating share profitability was the same for Experiment 2.

References

- Bröder, A. (2000). Assessing the empirical validity of the “Take-The-Best” heuristic as a model of human probabilistic inference. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *26*, 1332–1346.
- Bröder, A. (2002). *Decision making with the adaptive toolbox: influence of environmental structure, personality, intelligence, and working memory load*. Manuscript submitted for publication.
- Chater, N., Oaksford, M., Nakisa, R., & Redington, M. (in press). Fast, frugal and rational: How rational norms explain behavior. *Organizational Behavior and Human Decision Processes*.
- Czerlinski, J., Gigerenzer, G., & Goldstein, D. G. (1999). How good are simple heuristics? In G. Gigerenzer, P. M. Todd, & The ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 97–118). New York: Oxford University Press.
- Dhami, M.-K., & Ayton, P. (2001). Bailing and jailing the fast and frugal way. *Journal of Behavioral Decision Making*, *14*, 141–168.
- Elwyn, G., Edwards, A., Eccles, M., & Rovner, D. (2001). Decision analysis in patient care. *The Lancet*, *358*, 571–574.
- Ford, J. K., Schmitt, N., Schectman, S. L., Hults, B. M., & Doherty, M. L. (1989). Process tracing methods: Contributions, problems and neglected research questions. *Organizational Behavior and Human Decision Processes*, *43*, 75–117.
- Gigerenzer, G. (2000). *Adaptive thinking: Rationality in the real world*. Oxford: Oxford University Press.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, *103*, 650–669.
- Gigerenzer, G., Hoffrage, U., & Kleinbolting, H. (1991). Probabilistic mental models: A Brunswikian theory of confidence. *Psychological Review*, *98*, 506–528.
- Gigerenzer, G., & Selten, R. (2001). Rethinking rationality. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 1–13). Cambridge, MA: MIT Press.
- Gigerenzer, G., & Todd, P. M. (1999). Fast and frugal heuristics: The adaptive toolbox. In G. Gigerenzer, P. M. Todd, & The ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 3–34). Oxford: Oxford University Press.
- Gigerenzer, G., Todd, P. M., & The ABC Research Group (1999). *Simple heuristics that make us smart*. Oxford: Oxford University Press.
- Goldstein, D. G., & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, *109*, 75–90.
- Goldstein, D. A., Gigerenzer, G., Hogarth, R. M., Kacelnik, A., Kareev, Y., Klein, G., Martignon, L., Payne, J. W., & Schlag, K. H. (2001). Group report: Why and when do simple heuristics work?. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 173–190). Cambridge, MA: MIT Press.
- Harvey, N., & Bolger, F. (2001). Collecting information: Optimizing outcomes, screening options, or facilitating discrimination. *Quarterly Journal of Experimental Psychology*, *54A*, 269–301.
- Hastie, R., & Park, B. (1986). The relationship between memory and judgment depends on whether the judgment task is memory based or on-line. *Psychological Review*, *93*, 258–268.
- Jones, S., Juslin, P., Olsson, H., & Winman, A. (2000). *Algorithm, heuristic or exemplar: Process and representation in multiple cue judgment*. Paper presented at the Cognitive Science Society, Philadelphia.
- Juslin, P., & Persson, M. (2002). PROBABILITIES from EXEMPLARS (PROBEX): A ‘lazy’ algorithm for probabilistic inference from generic knowledge. *Cognitive Science*, *95*, 1–45.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–291.
- Martignon, L., & Hoffrage, U. (1999). Why does one-reason decision making work? A case study in ecological rationality. In G. Gigerenzer, P. M. Todd, & The ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 119–140). New York: Oxford University Press.
- Montgomery, H. (1983). Decision rules and the search for a dominance structure: Towards a process model of decision making. In P. C. Humphreys, O. Svenson, & A. Vari (Eds.), *Analyzing and aiding decision processes* (pp. 343–369). Amsterdam: North-Holland.
- Montgomery, & Willén, (1999). Decision making and action: The search for good structure. In P. Juslin & H. Montgomery (Eds.), *Judgment and decision making: Neo-Brunswikian and process-tracing approaches*. Mahwah, NJ: LEA.
- Newell, B. R., Rakow, T., Weston, N. J., & Shanks, D. R. (2002). *Search strategies in decision-making: the success of ‘success’*. Manuscript submitted for publication.
- Newell, B. R., & Shanks, D. R. (2003). Take-the-best or look at the rest? Factors influencing ‘one-reason’ decision making. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *29*, 53–65.
- Nygren, T. E., & White, R. J. (2001, November). *Effects of analytical and intuitive decision making styles on betting behavior in horse race gambling*. Paper presented at the 42nd annual meeting of the Psychonomic Society, Orlando, FL.
- Payne, J. W. (1976). Task complexity and contingent processing in decision making: An information search and protocol analysis. *Organizational Behavior and Human Performance*, *16*, 366–387.
- Rieskamp, J., & Hoffrage, U. (1999). When do people use simple heuristics and how can we tell? In G. Gigerenzer, P. M. Todd, & The ABC Research Group (Eds.), *Simple heuristics that make us smart* (pp. 141–167). Oxford: Oxford University Press.
- Seeley, T. D. (2001). Decision making in superorganisms: How collective wisdom arises from the poorly informed masses?. In G. Gigerenzer & R. Selten (Eds.), *Bounded rationality: The adaptive toolbox* (pp. 249–262). Cambridge, MA: MIT Press.
- Siegel, S., & Castellan, N. J. (1988). *Nonparametric statistics for the behavioral sciences* (2nd ed). New York: McGraw-Hill.
- Simon, H. A. (1956). Rational choice and the structure of environments. *Psychological Review*, *63*, 129–138.
- Slegers, D. W., Brake, G. L., & Doherty, M. E. (2000). Probabilistic mental models with continuous predictors. *Organizational Behavior and Human Decision Processes*, *81*, 98–114.
- Svenson, O. (1992). Differentiation and consolidation theory of human decision making: A frame of reference for the study of pre- and post-decision processes. *Acta Psychologica*, *80*, 143–168.
- Svenson, O. (1996). Decision making and the search for fundamental psychological regularities: What can be learned from a process perspective? *Organizational Behavior and Human Decision Processes*, *65*, 252–267.

Timmermans, D. (1993). The impact of task complexity on information use in multi-attribute decision making. *Journal of Behavioral Decision Making*, 6, 95–111.

Wiggins, J. S. (Ed.). (1996). *The five factor model of personality: Theoretical perspectives*. New York: The Guilford Press.

Received 14 March 2002