

RUNNING HEAD: FRAMEWORKS FOR SEARCH

Evaluating three frameworks for the value of information:
Adaptation to task characteristics and probabilistic structure.

Tim Rakow¹, Ben R. Newell², Kathryn Fayers¹, Mette Hersby¹

¹University of Essex, Colchester, UK

²University of New South Wales, Sydney, Australia.

Address for correspondence:

Tim Rakow

Department of Psychology

University of Essex

Wivenhoe Park

Colchester

Essex

CO4 3SQ

UK

Email: timrakow@essex.ac.uk

Tel: +44 (0) 1206 873812

Fax: +44 (0) 1206 873801

Evaluating three frameworks for the value of information:

Adaptation to task characteristics and probabilistic structure.

Abstract

We identify, and provide an integration of, three frameworks for measuring the informativeness of cues in a multiple-cue judgment task. Cues can be ranked by information value according to expected information gain (Bayesian framework), cue-outcome correlation (Correlational framework), or ecological validity (Ecological framework). In three experiments, all frameworks significantly predicted information acquisition, with the Correlational (then the Bayesian) framework being most successful. Additionally, participants adapted successfully to task characteristics (cue cost, time pressure, and information limitations) – altering the gross amount of information acquired, but not responding to more subtle features of the cues' information value that would have been beneficial. Rational analyses of our task environments indicate that participants' behavior can be considered successful from a boundedly rational standpoint.

Introduction

Consider an employee of a motor insurance company who is charged with the task of predicting which of two clients will make larger value claims in future years. Such predictions will often be wrong, but there are cues (age, sex, etc.) that may allow him to perform better than chance at this task (see Table 1). However, none of these cues perfectly predicts which client will have the greater claim value. Performing well in this task might require one to learn several probabilistic cue-criterion relations – hence, tasks of this kind are often described as a multiple-cue probability-learning (MCPL) tasks (Smedslund, 1955).

TABLE 1 ABOUT HERE

The MCPL paradigm has fostered hundreds of experimental studies (Holzworth, 2001), though typically the focus is on judgments for individual objects rather than choice between multiple objects. Many ‘everyday’ choice tasks share properties of the MCPL choice task outlined above. When choosing between applicants, an employer needs to identify who will perform best in the job on the basis of cues that predict only imperfectly (prior experience, qualifications, etc.). When identifying which of two patients in the emergency room is more seriously ill, the doctor may consider whether each of the patients is conscious, short of breath, complaining of pain, and so on. Often the relevant cues are not immediately apparent – in which case, information search is presumably neither random nor exhaustive. Pre-decisional information search is the focus of this paper. Specifically, what determines which cues are considered and the frequency with which they are examined over repeated choices?

To address these questions we focus on two types of feature of the MCPL domain: (1) the constraints or demands placed upon the decision maker by task characteristics, and (2) the cue-criterion relationships for the probabilistic environment in which the decision maker operates. The latter set of features rarely lie within human control, the former set may do so

more frequently. We therefore focus on two different aspects of the decision makers' adaptivity – their ability to alter search behavior in response to changing task characteristics, and their ability to learn features of the environment that can usefully guide information search.

Adapting to Changing Task Characteristics – Stopping Search and Deciding

The most commonly examined task characteristics in relation to information acquisition are information costs and time pressure. With respect to information costs, the standard summary is that participants do adapt appropriately to differences in cue costs, but do so insufficiently. Relative to optimal purchasing, participants buy too much information when little is needed, and under-purchase when higher levels are beneficial (Connolly & Serre, 1984; Rapoport & Wallsten, 1972). Much of this work was conducted under the Bayesian optional stopping paradigm (Edwards, 1965), where a participant can purchase samples from a population. Conceptually, this is a little different to the multiple-cue paradigm that we consider – it corresponds to repeated sampling of a single cue rather than sampling multiple cues. We examine adaptation to cue cost in Experiment 1.

Not surprisingly, under time-pressure people do search for less information prior to making a choice, though reductions in information acquisition can be less than might be anticipated as people can also adapt by speeding up their processing of information (Payne, Bettman, & Johnson, 1993). Payne et al. (1993) observed a tendency to employ attribute-based strategies in preference to alternative-based strategies under time pressure. Such strategies allow one to eliminate options on the basis of partial information, obviating the need to fully evaluate options on the basis of all information. We examine adaptation to time pressure in Experiment 2.

The final task characteristic we consider, in Experiment 3, is changes in the scope of information search permitted. Specifically, we examine the effect upon search patterns of changing the amount of information that can be acquired. This corresponds to fixing the budget for pre-decisional information (for a given decision) – and, across decisions, considering the effects of different budget levels. Additionally, we compare situations in which there is certainty about the number of pieces of information available to make a decision before a given trial commences, with those on which there is uncertainty. Thus we compare perfect knowledge of the limits upon information acquisition with the imperfect knowledge that prevails in many competitive markets, where the exact time at which a decision must be made is not predetermined. In such situations, one does not know in advance how much information can be acquired before being called upon to make a decision (because one does not know how much time will be allowed for information acquisition).

In all three experiments we expect adaptation - primarily at the level of the amount of information acquired - to these changes in task characteristics. However, we argue that there are more subtle aspects of the environments that participants could adapt to – specifically, some particular features of the probabilistic structure of the environment and how these affect the ‘value’ of information.

Adapting to features of the probabilistic environment – search patterns and the value of information

The standard prescription for pre-decisional information search when multiple cues (or ‘attributes’) are relevant to the choice at hand is to search cues in descending order of importance (Albert, Aschenbrenner, & Schmalhofer, 1989), or at least to do so probabilistically (Tversky, 1972). Where choice is not purely a matter of satisfying idiosyncratic personal preferences but demands maximising value according to an objective

criterion, how should cue importance or the value of information be defined? We have been able to identify three different approaches that provide slightly differing answers to this question – we term these the Bayesian, the Correlational, and the Ecological approaches. Each approach defines the magnitude or strength of the cue-criterion relationship in a particular way – thus these different approaches sometimes differ in how they rank order cues by importance.

The Bayesian approach. The Bayesian framework lends itself to the consideration of information search, and the relative value of different pieces of information. Bayes Theorem can be used to update the subjective probability of an event or the confidence in some proposition upon the receipt of additional information. It is therefore quite natural to conceive of the Bayesian decision maker as one who revises belief step-by-step upon the sequential acquisition of cue information – and chooses between options when sufficiently certain of the best option. Information is then valued according to the magnitude of the reduction of uncertainty or the revision of belief that it yields. Because we cannot know with certainty exactly how to revise belief until we have acquired the relevant information, the appropriate Bayesian measure for determining the order of information acquisition is the *expected change in belief* (Lindley, 1956). Expected Change in Belief (ECB), also known as *expected information gain* (Oaksford & Chater, 1998), has been used to account for human information search strategies consistent with a positive test strategy that have been observed in several cognitive tasks that require information acquisition (Klayman & Ha, 1987). In the context of MCPL choice tasks, the Bayesian approach would prescribe that at a given stage in the search process the next cue to be acquired should be the cue that yields the greatest ECB (if all cues are equally costly).

The Correlational approach. Natural, behavioral and social scientists are quite used to measuring the strength of a relationship according to correlation coefficients, and to ranking

the strength of different relationships accordingly (e.g., as typified by the use of Pearson's r as a measure of *effect size*, Cohen, 1969). In the context of judgment and decision making under uncertainty, Egon Brunswik proposed that successful judgment of an external criterion variable in probabilistic environments depended upon weighting cues according to their strength of relationship with the criterion variable (Brunswik, 1955). Thus judgment and choice involving objective criteria should mimic cue-criterion relations that are observable in the environment. Social judgment theorists who take their lead from Brunswik have usually (though not exclusively) assessed these relations using correlational measures such as Pearson's r or beta-weights derived using multiple linear regression (Hammond & Stewart, 2001; Slovic & Lichtstein, 1971). Whilst what Brunswik provided was more a theory of information integration and less a theory of information search, he was explicit that not all cues would be examined on all occasions that a particular multiple-cue judgment is made, as, often, cues can effectively and efficiently be substituted for one another ('vicarious functioning'). Thus it is possible to conceive of cue weights as the value ascribed to a cue in the integration process, and/or the propensity to attend to (or search for) a cue. Empirical investigations of information purchase indicate more frequent acquisition of cues that are more strongly correlated with a criterion, consistent with partial (though imperfect) sensitivity to variation in cue-criterion correlation (Connolly & Serre, 1984).

The Ecological approach. Gigerenzer and colleagues have forwarded a vision of human judgment and choice built upon an 'adaptive toolbox' consisting of 'simple heuristics' (Gigerenzer & Todd, 1999; Todd & Gigerenzer, 2000). These relatively uncomplicated strategies search through information in an order determined by the cue-criterion environment, then stop search and choose (or estimate) according to clearly defined rules. One of these simple heuristics, Take The Best (TTB), is applicable to the kind of MCPL choice tasks that we consider in this paper (Gigerenzer & Goldstein, 1996). TTB searches

through cues in descending order of their *ecological validity* - where ecological validity (EV) is defined as the probability of a correct choice *given* that the cue points to one object. The reference class for this probability is the set of all possible objects pairs in the consideration set. In addition to its EV, a cue's *discrimination rate* is also of relevance to the value of the information 'carried' by the cue. The discrimination rate (DR) is the probability that a cue will point to a unique choice (again defined for all possible object pairs). For instance, in the case shown in Table 1, the first cue (sex) does discriminate between the objects - and is likely to do so relatively frequently for this task, probably yielding a DR close to the maximum possible value of 0.5. The third cue (criminal record) does not discriminate in the case shown, and in general, rarely will (as for most pairs of clients *neither* will be a convicted felon) – resulting in a DR close to the minimum possible value of zero. However, of these two cues, criminal record may well have a higher EV. If so, when it can be used to make a choice, it is the more predictive of the two cues. It will be seen that searching through cues by descending EV places a premium on accuracy, with the potential drawback that many cues must be searched through before a discriminating cue is examined.

The relationship between these approaches. It can be shown that the Bayesian and Correlational measures of the cue-criterion relationship can be represented as functions of both the EV and DR of cues. Thus the Ecological approach provides a scheme for decomposing ECB and Pearson's correlation into components that depend upon the EV and DR of cues.

Working within the ecological framework, Martignon and Hoffrage (1999) defined an additional measure of the cue-criterion relationship within a MCPL choice environment that they termed *success*. Success is the probability of a correct choice when only that cue is used, and in the context of two-alternative forced-choice (2AFC) success (S) is defined as:

$$S = EV*DR + (1 - DR)*0.5 \quad (1a)$$

or equivalently, by rearrangement, as:

$$S = (EV - 0.5)*DR + 0.5 \tag{1b}$$

We have demonstrated in another paper that ranking cues according to ECB is isomorphic to ranking them by success (Rakow, Hinvest, Jackson, & Palmer, 2004). It should make intuitive sense that how much one expects to revise belief on examination of a cue depends upon the probability that one will revise at all (dependent upon DR) and the magnitude of revision given that one does revise one’s belief (dependent upon EV). Thus ECB can be seen to represent the value of information as a mix of DR and EV.

Similarly, it can be shown that the correlation between the pattern of cue values and outcome can also be represented as a mix of DR and EV. To do so requires representing the pattern of cue values as a single variable with three (ordered) levels: ‘+1’ if the pair of cue values point to object A, ‘0’ if the cue does not discriminate, and ‘-1’ if the pair of cue values point to object B. Then the Pearson correlation coefficient for the relationship between cue pattern and the choice criterion (A is greater vs B is greater) can be represented as:

$$\text{Pearson's } r = (2*EV - 1)* (DR) \tag{2}$$

A demonstration of this result is given in Appendix 1. Again it should make some intuitive sense that this correlation coefficient should depend upon both EV and DR. The EV reflects co-occurrences of discriminating cue patterns and choice criterion values – and the frequency of these co-occurrences influences the strength of correlation. The DR reflects the *variance* in cue patterns – which affects the strength of correlation (as correlation is a standardized measure of *covariance*).

Comparison of equations 1b and 2 indicate that the ranking of cues according to success/ECB can be different to that defined by Pearson’s *r* (if EV and DR vary across cues). Both sets of rankings reflect the DR and EV of cues, but of the two sets of rankings Pearson’s *r* will be more influenced by variation in EV. (i.e., If cues are selected in descending order of

Pearson's r , this order will be more similar to the descending order of EV than would be the case if cues are selected in descending order of success/ECB. For instance, in the first two experiments that we report, it is sufficient to switch the order of one pair of (the four) cues to rearrange the order of cues from descending order of Pearson's r to the descending order of EV. Effecting the equivalent rearrangement from success/ECB order into EV order requires two pair-wise reversals.)

We have previously found that success/ECB provides a better description of participants' cue search than EV or DR alone (Newell, Rakow, Weston, & Shanks, 2004). However, for the probabilistic environment we employed in those previous experiments the rank order of cues by success/ECB was identical to that of Pearson's r , thus preventing us from assessing the ability of the three frameworks for describing behavior. Furthermore, Newell et al. (2004) did not manipulate any aspect of the task environment - cue costs were kept constant, and participants always operated without the imposition of any time pressure. In the experiments that follow we created environments in which *all three* orders (prescribed by the Ecological, Correlational and Bayesian approaches) are different, and investigate how well each of the three approaches describe search behavior under different conditions of time, information cost and scope of information acquisition.

Inferring importance from observed patterns of search. Perhaps the most intuitive method for inferring search behavior from the data is to examine the extent to which participants look through cues in the strict deterministic orders predicted by the different conceptions of importance. However, using such deterministic orderings as a benchmark precludes the possibility that external factors may affect participants' search. For example, Albert et al. (1989) suggested that though important cues should be processed earlier than less important ones, in many settings this ordering may not be perfect as other factors such as

reading habits, accessibility, salience and availability could all play a role. Similarly, Tversky (1972) proposed a *probabilistic* search rule for the elimination-by-aspects model of choice (a model that is closely related to the lexicographic strategies we consider). In the model, cues (or ‘aspects’) are selected with a probability proportional to their weight or importance for the particular choice at a particular moment in time. As Tversky (1972) stated, “the probabilities merely reflect the fact that at different moments in time different states of mind (leading to different choices) may prevail” (p. 296).

In the following experiments we use a computerized information board set up in which it is quite likely that external factors – such as ease of moving the mouse to particular buttons – will have an effect on the order of search. Consider a participant who decides to examine the two most ecologically valid cues first but due to their position on the screen always looks at the second most valid followed by the most valid. Such behavior would be inconsistent with a strict validity ordering, yet would seem to indicate that the participant was responding primarily to EV.

For these reasons we examined the *frequency* with which particular cues were acquired. That is the proportion of trials on which each cue was acquired aggregated across participants. We contend that an ordering based on aggregated frequency of acquisition provides a valid indication of the extent to which particular cues are considered ‘important’ by participants.

Experiment 1: Adaptation to the cost of information

Consistent with prior investigation of information purchase, we anticipate attendant changes in the number of cues bought prior to choice as cue-cost increases. However, our discussion of search frameworks demonstrates that the cue-criterion relationship can be decomposed into components that depend upon the DR and EV of cues. We can therefore

examine whether changes in the patterns of cue purchase will reflect these dimensions and mediate changes in the amount of information bought. For instance, when information becomes more expensive, will participants simply truncate their search, or will they search for cues with higher DRs (those more likely to discriminate between options) at the expense of cues with higher EVs (those more likely to indicate the ‘correct’ option)? To illustrate the rationality of these adaptations, we present a rational analysis of the probabilistic environment employed in Experiment 1.

Rational Analysis of the Environment for Experiment 1

The task in Experiment 1 was to predict, using up to four binary cues, which of two objects would have the higher criterion value. Table 2 summarizes the cue-criterion relationships according to the measures that could define cue importance. (See method section for details of the design that generated these relationships, and for a description of the ‘share prediction’ cover story given to participants.) We considered five classes of search strategy, which previous work had suggested would yield differing levels of accuracy and use differing amounts of information. (1) A (compensatory) *weighted additive rule* (WADD) derived using multiple linear regression – this strategy always searches through all cues. The next three classes of strategy all use one-reason stopping and decision rules (‘search until the one cue discriminates, and choose the object that the cue points to’), but search through cues according to different hierarchies of importance. (2) The *Ecological strategy* searches cues in descending order of EV, and is therefore the TTB strategy described by Gigerenzer and Goldstein (1996). We consider both the ‘Full’ strategy, which searches through as many cues as are available before guessing (if no cues discriminate), as well as ‘truncated’ strategies that stop and guess after a maximum of one, two, or three cues if none have discriminated by that point (Ecological-T1, Ecological-T2, and Ecological-T3, respectively). (3) The *Bayesian*

strategy searches through cues in descending order of success/ECB – and, again we consider the entire family of full and truncated strategies (-Full, -T3, -T2, and -T1). (4) The *Correlational strategy* searches cues in descending order of Pearson's r – and, similarly, we consider the entire family of Correlational strategies. (5) *Random Choice* (RC) searches for no information and chooses one of the two objects randomly (and therefore expects to make the correct choice on 50% of occasions). This provides a baseline against which other classes of strategy can be assessed.

TABLE 2 ABOUT HERE

FIGURE 1 ABOUT HERE

Expected payoffs for each strategy were computed, setting the reward for a correct choice at 100 points, and allowing the cost (in points) of examining each cue to vary. Figure 1 shows that the strategy yielding the greatest expected payoff varies as the cue costs increase – reflecting the trade-off between the accuracy achieved and the amount of information used by these strategies. Each intercept on the vertical axis shows the expected accuracy (% of correct choices) of that strategy. The gradient of each line reflects the amount of information searched for, with more frugal strategies having more shallow gradients indicating that payoffs for these strategies are less sensitive to changes in the cost of information. For a given cost (corresponding to a point on the horizontal axis), the best strategy from among those considered can be found by observing which line is uppermost. For clarity, only those strategies that yield the highest expected pay-off over some range of costs (however short) are shown in Figure 1 – i.e., dominated strategies are not plotted. Naturally, when information is free the most accurate strategy (WADD) provides the highest expected payoff. However, even with minimal search costs, the more frugal (and only slightly less accurate) Ecological-Full strategy returns a higher expected payoff. When search costs are moderate, the Correlational and then Bayesian strategies yield the greatest rewards, though the

maximum number of cues one is prepared to search should decrease with increasing costs.

When search costs are high, no cues are sufficiently informative to improve upon the expected pay-off from random choice. Thus, as information costs increase two types of adaptation to search are seen to be of benefit. Firstly, changing the order in which cues are searched – switching from orders more akin to EV, to orders that are more influenced by the DR of cues (i.e., switching progressively from the Ecological to the Correlational and then to the Bayesian search orders). Secondly – reducing the maximum number of cues that one is prepared to search through prior to making a choice.

Taking data from previous similar experiments into account (Newell et al., 2004), this rational analysis was used to guide the cost of information that was set in four different cost conditions (3, 10, 17, and 24 points per cue). In these earlier experiments, participants differed in the amount of information they purchased, and, on average, tended to over-purchase information relative to what would be expected given *purely* financial considerations (Newell et al., 2004). For instance, with cue costs set at a level equivalent to 17 points, just over half of the participants examined more than two cues per trial (on average). Consequently, it was felt necessary to set costs up to 24-point per cue if we were to examine search behavior in situations where participants opt to restrict their search significantly.

There is no reason to preclude the validity of non-financial motivations in experiments such as this (e.g., some intrinsic reward for accuracy), and that such motivations might legitimately vary between individuals. Therefore, we cannot specify that the strategy our analysis indicates is ‘best’ for a given cost condition is necessarily the one that *should* be adopted by all participants. However, over a range of cost conditions (such as those set) we might expect to see adaptations in search behavior consistent with some, or all, of the sequence of strategy changes evident in the rational analysis above. Thus, with a more

qualitative interpretation of our rational analysis we would anticipate that as information costs increase there would be a tendency to switch or adapt strategies. Switching from compensatory decision making to strategies that are prepared to guess even if additional information could help improve accuracy. In addition, a qualitative interpretation of our analysis would anticipate an EV-based (Ecological) search rule with lower information costs, but a Pearson's r (Correlational) or ECB/success-based (Bayesian) search rule with higher costs. Thus, our rational analysis shows that adaptations in *both* the stopping rule *and* the search rule can be used to enhance the pay-off from choices.

Method

Participants

Twelve members of the University College London community took part in the experiment. Six were male and six were female. They had a mean age of 21.9 years (range 18 – 33, SD = 4.3). Each participant completed four separate testing sessions.

Stimuli and Design, and Procedure

A share profitability prediction task similar to the one used in a number of previous studies was used (Bröder, 2000; Newell et al., 2004; Newell & Shanks, 2003; Newell, Weston, & Shanks, 2003). 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 described by four binary cues with semantic labels concerning aspects of the company's financial status (*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?*). Each cue had an associated validity, discrimination rate, success rate and Pearson's r .

Repeated choices were presented to participants in blocks of 64 trials. Each block was generated (independently) according to a fractional factorial design using SPSS ORTHOPLAN, to ensure that all four cues were independent (i.e., uncorrelated). Four binary cues were created such that $1/8$, $1/4$, $3/8$ or $1/2$ of the generated sets of cue-values were positive, resulting in cue discrimination rates of 0.22, 0.37, 0.47 and 0.50 respectively. Twelve blocks were created in total and these were rotated such that each block was used an equal number of times in the training, practice and test phases of the experiment (see below).

Table 2 displays the validities, discrimination rates, success rates and Pearson's r of the four cues A-D. The assignment of these properties to the four nominal (binary) cues was counterbalanced across participants, but the screen position of the cue labels was constant for all participants. By counterbalancing the position of cues, we remove any systematic effect of the location in which cues appear on the screen. By counterbalancing cues across the four labels, any systematic effect of prior beliefs about the supposed differential value of cues (that participants might infer from the labels) is also eliminated.

For each comparison there was an associated probability of Share A being the most profitable. After each choice, the probability that the chosen share was most profitable was computed according to Bayes' Rule, assuming stochastic independence of the cues (details of the calculations can be found in Newell & Shanks (2003)). A random number generator then determined which share became most profitable according to this probability.

Each testing session was divided into three stages: a training block, a practice block and a test block. In the training block, consisting of 64 trials, cue information was supplied automatically, simultaneously, and at no cost to participants. On each trial, participants were motivated to learn about the cues by winning 20 points for every correct prediction. Participants were told that every 500 points was equal to UK£1.00 (approx. \$1.65; Eur 1.43) and that their points would be converted into money at the end of the fourth and final testing

session. Following training, participants were given a practice block in which cue information had to be bought. Participants were able to buy information about each cue in any order and were required to buy at least one cue. This feature was included to ensure that data on cue-search (our principal interest) was obtained for all participants. Participants bought information by clicking on a screen-button that uncovered the ‘values’ of each piece of information for the two shares (e.g., a “YES” for Share A and a “NO” for Share B in answer to the question “Is it an established company?”).

The cost of each piece of information varied across testing sessions (see below), but the pay-off for a correct prediction was held constant at 100 points. Participants were told that any points spent or won during the practice block *would not* affect their final earnings. They were encouraged to use the 64 trials of the practice block to work out how to earn the most money in the test block that followed practice. The information cost and pay-off structure of the 64 trials in the test block were the same as those in the practice block but points won *were converted into final earnings*. The trials in each block were presented in an order randomly selected for each participant.

Information cost conditions. Our principal interest was to examine whether the relative cost of information would affect not only the *amount* of information purchased but also the frequency with which different cues were purchased. To this end we examined conditions in which the cost of information was 3, 10, 17 or 24 points per cue (see the rational analysis for an explanation of why these particular values were chosen). In a single testing session the cost of information was held constant (i.e., if information cost 3 points in the practice block, it cost 3 points in the test block). The order in which participants were run in the four cost conditions was counterbalanced. Testing sessions were run either on four separate days or over two days with a session in the morning and a session in the afternoon.

Procedure. Participants were given a brief verbal description of the experiment telling them they would be making decisions about the profitability of company shares. Participants were then given full written instructions on the screen, which they read through with the experimenter, and on a handout placed on the table next to the computer. Participants were told that they had to complete four testing sessions (over the course of two or four days) and that the cost of information would be different in the different sessions. Participants were also 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. Outcome feedback was provided after each prediction during all stages of the experiment (training, practice and test). (A figure depicting the screen layout and a more detailed description of the procedure can be found in Newell, et al. (2004)).

At the end of the fourth testing session participants were thanked, debriefed, and paid their total earnings.

Results and Discussion

Purchase Proportions

The frequency with which each cue was purchased during the test blocks was recorded for each participant in each cost condition (summarized in Table 3a). Table 3a indicates that the rank order of purchase proportions fell in the order prescribed by Pearson's r in all cost conditions. A two-way (cost condition by cue) repeated measures ANOVA was performed with the purchase proportion as the dependent variable. A significant main effect of cost, $F(3,33) = 8.08$, $p < .001$, and a significant linear trend, $F(1,11) = 13.88$, $p = .003$, indicated that participants did acquire less information when it was more expensive.¹

INSERT TABLE 3 ABOUT HERE

Tukey HSD post hoc tests indicated that purchase levels in the 3-point condition were significantly higher than those in the 24-point ($p < .001$) and 17-point ($p < .01$) conditions,

and that more information was purchased in the 10-point condition than in the 24-point condition ($p < .05$). Thus, purchase levels in adjacent cost conditions did not differ significantly from each other, whilst all pairs of non-adjacent cost conditions did differ in the expected direction. There was a significant main effect of cue, $F(3,33) = 6.15$, $p = .002$.

Tukey HSD post hoc tests indicated that cue B (lowest EV, success/ECB and Pearson's r) was bought significantly less often than cue D (highest success/ECB and Pearson's r) and cue A (greatest EV), with $p < .01$ and $p < .05$, respectively. There was no significant two-way interaction, $F(9,99) = 0.83$, $p = .595$. Patterns of cue purchase did not differ significantly by cost condition, and changes in purchase frequency across different cost conditions did not differ significantly from one cue to another. Thus there is clear evidence that people search for less information when costs are higher, but no evidence that changes in cost prompt them to change the relative frequency with which they purchase particular cues.

A simple tally of which cue was examined most frequently by each of the participants showed that cue A (highest EV) and cue D (highest success/ECB, and highest Pearson's r) were the 'most favored' cue for five participants each. Cues B and C were examined with greatest frequency by only one participant each.

Purchase proportions were further examined using Page's non-parametric trend test to determine which of the four specified cue hierarchies (by EV, DR, success/ECB, or Pearson's r) described participants' search patterns best. Purchase proportions were positively related to each of the four trend orders examined (i.e., $z > 0$ in each case), and always significantly related to the orders predicted by EV, success/ECB, and Pearson's r (see Table 3b). Purchase proportions were always best predicted by the order determined by Pearson's r . In three of the four cost conditions, purchase proportions were better predicted by success/ECB order than by EV order, with the reverse being true in the 17 point cost condition.

Stopping Search

We examined two types of departure from the one-reason stopping rule ('stop search at the first cue that discriminates'): (1) the purchase of 'unnecessary information' – obtaining further cues after the first cue that discriminates, and (2) 'Unforced guessing' – choosing a share when no cue has discriminated *and* not all cues have been examined. Figure 2 displays these two types of departure. Taking both types together, there are departures from the one-reason stopping rule on half of the test trials in the highest cost condition, and about one third of the test trials in the other three conditions. There is a significant effect of cost condition on the purchase of unnecessary information ($F(3, 33) = 3.87, p = .031$) reflecting the lower levels of purchase in higher cost conditions (though the linear trend did not reach significance ($F(1,11) = 4.05, p = .069$)). As can be seen in Figure 2, unnecessary information is rarely purchased in the two highest cost conditions. The proportion of unforced guessing increases steadily with information cost, with a significant linear trend ($F(1,11) = 15.67, p = .002$).

Taken together, the frequency and stopping rule analyses indicate that people adapt (appropriately) to higher information costs by truncating their sequential search through cues rather than by changing the frequency of purchase of particular cues to a strategy that is more likely to yield discriminating information. We consider the reasons for this behavior in the General Discussion.

FIGURE 2 ABOUT HERE

Experiment 2: Adaptation to time constraints

Financial costs are not the only information that people have to contend with. Arguably, in every day situations it is more common for people to experience information costs in terms of the time taken or effort expended in searching for information. We investigate the effect of time pressure by examining a fixed number of choices made within an overall time limit for a

set of choices. We manipulate search time by altering the length of time participants must wait after acquiring a cue before they may acquire the next cue. Cues are acquired without financial cost, but correct decisions are rewarded as before, and choices not made within the overall time limit incur a financial penalty. Thus, this experiment examines choice when ‘time is of the essence’ – those difficult situations where additional information makes the right choice more likely, but the time taken to acquire information risks costly delays in reaching a decision. This is the dilemma faced by the emergency medicine doctor or the financial trader. In examining a patient for additional signs and symptoms, the doctor risks delaying necessary treatment (for this and for other patients). By researching the wisdom of an investment, the trader risks missing the opportunity to invest at the price that was initially quoted, or may be indisposed when other investment opportunities arise.

Time and effort in acquisition can genuinely be considered to be search costs (Fried & Peterson, 1969; Payne, 1982), however, this time-limited paradigm differs in one important respect from the financial cost paradigm used in Experiment 1 and previously (e.g., Newell et al., 2004). In Experiment 1, the cue cost function was linear – each additional cue acquired incurred the same incremental cost (within a given cost condition). When working within a time limit, the cue cost function is a positively increasing one. If the first few cues acquired can be examined sufficiently quickly that there is little danger of failing to keep to the time limit, these can be considered to have relatively little cost. If acquiring further cues jeopardizes one’s ability to make a sufficient number of choices within the time limit, these are high cost cues.

In addition to anticipating the truncation of information search under high time pressure, we can again examine whether people adapt their patterns of cue search with respect to the mix of DR and EV and determine which cues are examined with relatively greater frequency. Specifically, under high time pressure, will cues of high DR (those more

likely to discriminate between options) be examined relatively more frequently than under low time pressure? If so, we would expect a shift from search strategies in keeping with the Ecological framework under low time pressure, but following the prescriptions of the Bayesian and Correlational frameworks under high time pressure.

Method

Participants

Thirty-two members of the University of Essex community took part in the experiment. Ten were male and 22 were female. They had a mean age of 25.5 years (range 17 – 40, SD = 5.3).

Stimuli and Design, and Procedure

The share prediction task as described in Experiment 1 was employed. Cue EVs (and consequently success/ECB and Pearson's r) were altered slightly from Experiment 1 such that only on those trials with no discriminating cues was the probability of both shares being the better share equal (at 0.5). This ensured that there was always one choice that was *strictly* better than the other. The relative ranking of cues with respect to the parameters remained as it had been in Experiment 1 (e.g., the cue with highest EV had the second highest Pearson's r and the third highest success/ECB in either experiment). Cues were placed (randomly) into a new cyclical sequence (see Table 4).

TABLE 4 ABOUT HERE

Participants again earned points on each trial, according to the payment function: $100(p - 0.5)$, where p is the probability of a correct choice. Thus a participant performing at chance would earn no points. Points were incremented to a running total after each trial.²

Time pressure was manipulated within-participants with each participant encountering three blocks of 64 trials: a training block, a low time-pressure block, and a high time-pressure

block. In the training block, all cue information was provided simultaneously, without cost, and without any time limit. In the instructions given at the start of the experiment, participants were informed that they would later have to make decisions under time pressure and may not have time to examine all the available information - they were advised to use the training trials to learn which pieces of information were more informative. Trials in each time-pressure block were divided randomly into two fixed phases of 32 trials: a practice phase followed by a test phase. For each phase there was a time limit of 320 seconds, in which the 32 trials were to be completed. In other words, time pressure was applied at the level of the session, not at the level of the individual trial (participants having an *average* of 10 seconds -- 320 seconds /32 trials -- to complete each trial). At the end of every trial, information on screen was updated to show how many trials and how much time remained. On screen messages also indicated when one half and when three quarters of the time limit had elapsed. Each trial not completed within the limit incurred a penalty of 25 points. Points earned in the two test phases were converted into money at a rate of 400 points to UK£1.00, and were added to the participants 'turn-up' fee of UK£2.00.

In the Low time pressure condition, following the acquisition of a cue the program did not permit participants to acquire another cue for two seconds. In the High time pressure condition, this delay was programmed at three seconds. Thus, time pressure was manipulated (within subjects) by changing the length of time it took to acquire information, not the length of time in which choices had to be made.

Pilot work had indicated that these lengths in delay would result in a difference in the *average* amount of information that could be acquired in excess of half a cue per trial. A rational analysis of the environment had indicated that a difference of this magnitude was sufficient to alter which search strategy performs best. Obtaining the cues required for the Ecological-Full strategy (an average of 2.4 cues per trial) was estimated to be relatively easy

in the low time-pressure condition but relatively hard under high time pressure. In contrast, the Bayesian-Full strategy requiring an average search string of only 2.0 cues per trial could be utilized effectively even under high time pressure.

The screen position of cues was counterbalanced across four configurations as per Experiment 1, and this was fully crossed with the counterbalancing of the order of Time Pressure conditions. Consequently, four participants were assigned to each of the eight possible combinations of information position and condition order, eliminating any systematic effects of screen location (of cues), prior beliefs, and task learning.

Results and Discussion

Acquisition proportions

Table 5a summarizes the proportion of information examined and indicates that the rank order of acquisition proportions fell in the order prescribed by Pearson's r in both Time pressure conditions. A two-way (Time pressure by Cue) repeated measures ANOVA with proportion examined as the dependent variable revealed significant main effects of Time pressure, $F(1,31) = 8.60$, $p = .006$, and of Cue, $F(3,93) = 13.12$, $p < .001$, but no significant two-way interaction, $F(3,93) = 0.43$, $p = .731$. Tukey HSD post hoc tests indicated that cue A (highest success/ECB and greatest Pearson's r) was examined more frequently than each of the other cues ($p < .001$ for cues B and C, $p < .01$ for cue D). No other pairs of cues were significantly different. Thus, participants were sensitive to the time pressure manipulation, and did not examine cues with equal frequency.³

TABLE 5 ABOUT HERE

Tallies of each participant's 'most favored' cue found that 17.5 participants (55%) chose cue A (highest success/ECB and greatest Pearson's r) most frequently [.5 denotes two cues acquired with equal frequency]. Cue D (highest EV, and second highest Pearson's r)

was the most favored cue for 8.5 participants, with cues C and B being the most frequently chosen cues for only 3.5 and 2.5 participants respectively.

As shown in Table 5b, Page trend tests indicated that the frequencies with which each cue was examined are best predicted by the orders of importance prescribed by Pearson's r (Correlational framework) and success/ECB (Bayesian framework). Thus, there is clear evidence that the frequency with which a cue is examined is related to its 'importance' as defined by Correlational and Bayesian frameworks. There is also clear evidence that people examine less information under higher time pressure. However, there is no clear evidence favoring adaptive changes in the mix of DR and EV that people respond to as a function of time pressure.

Consistent with Experiment 1, adaptation to costs takes the form of acquiring less information, but this is not achieved by changing which cues are acquired *relatively* more, or less, frequently.

Experiment 3: Adaptation to uncertainty

Experiment 3 examines one final constraint upon information acquisition – *fixed resources for each individual decision*. This was achieved by placing a limit on the number of cues that could be acquired before a choice was made. Thus the constraint on information search applies to each individual trial, rather than being averaged across a block of trials (as was the case in Experiments 1 and 2). This models decisions that have a fixed time or fixed budget for information acquisition. Any choice with a strict deadline for selecting your option has this characteristic with respect to pre-decisional information acquisition.

On different trials in this experiment, participants could acquire one, two or three cues. In the absence of financial costs for cues or time limitations, we anticipated that

participants would acquire the amount of information that they were allowed to. However, *which cues* are acquired remains of interest.⁴

In addition to limiting the amount of information participants could acquire, we manipulated whether or not participants *knew in advance* how many cues they could acquire before making a decision. This models decisions where, after initiating the search for pre-decisional information, the deadline for selecting your option may be expedited (ceasing information acquisition) or delayed (allowing further search through cues). In one block of trials participants knew at the start of the trial exactly how many cues they could acquire before making a decision (Known Limit Condition). In another block participants discovered the limitation to their cue acquisition only at the point when a decision was required (Unknown Limit Condition). If participants are sensitive to the different informational properties of the cues, and adapt accordingly to these properties, then we expect a greater frequency of acquisition of higher EV, low DR cues when information limitations are known in advance. This is because when there is an opportunity to examine other cues, participants can ‘take the risk’ of acquiring a cue that has a lower probability of being useful for the decision (i.e., low DR). Such a risk is less attractive in the case of an unknown limit for acquisition because one might be left with only a non-discriminating cue when the choice must be made.

Method

Participants

Thirty-two members of the University of Essex community took part in the experiment. Ten were male and 22 were female, and they had a mean age of 22.8 years (range 16 – 46, SD = 6.6).

Stimuli and Design, and Procedure

The task and the measures obtained were the same as Experiment 2. A new environment was created for this experiment, with cues placed randomly into the cyclical order shown in Table 6. The DRs were slightly more widely and more evenly distributed – this was achieved by setting the proportion of positive cue values to $1/9$, $2/9$, $3/9$, and $4/9$ for the four cues. This configuration of cue values resulted in experimental blocks of 81 trials. In this new probabilistic environment, the cue with the highest EV had the lowest success (Cue B) – giving a clearer distinction between the orders prescribed by success/ECB and EV. Cue parameters (shown in Table 6) were such that for one or two cues examined, a Bayesian strategy was more accurate than an Ecological strategy (with absolute differences of 6% and 2%, respectively, calculated from rational analysis). In contrast, if three cues are examined the advantage for the Bayesian strategy disappears, and the Ecological strategy outperforms Bayesian strategy by a small margin (1%). This indicates that there is a small advantage to EV-based search when you know you can search for more information, with search directed by success/ECB being clearly advantageous when cue search is highly limited.

TABLE 6 ABOUT HERE

Each participant encountered three blocks of 81 trials: a training block, and the two test blocks that comprised the Known Limit and Unknown Limit Conditions. In the training block, all cue information was provided simultaneously. Participants were told that they would later have to make choices without having all the information available, and that they should use the training trials to learn which information would be most useful for this. Trials in each test block were divided into three levels of Amount of Information: participants could acquire one, two, or three cues before making their choice. Equal numbers (27) of each trial were experienced and the order of different types was assigned randomly but fixed across participants. In the Known limit condition, participants were told how many pieces of

information they could acquire via a message shown on the screen from the start of each trial. In the Unknown limit condition, when the maximum number of cues permitted on that trial had been acquired, the message “Choose now” was shown on the screen. The programming of the experiment did not permit participants to acquire more than the permitted number of cues. Points earned in the two test blocks were converted into money at a rate of 1000 points to UK£1.00, and were added to the participants’ ‘turn-up’ fee of UK£1.00.

As previously, the screen position of cues was counterbalanced across four configurations, and was fully crossed with the counterbalancing of the order of Known/Unknown limitation conditions.

Results and Discussion

Proportions of information examined - the effects of Cue and Certainty

To determine whether patterns of cue examination varied between Known and Unknown conditions we examined the effect of Cue (four cues, A to D) and Certainty (Known and Unknown) for each level of Amount of information (separate analyses for one, two, or three cues). All main effects of Certainty were non-significant ($p > .27$), as were all two-way interactions ($p > .15$). Hence, in later analyses data were collapsed across both Certainty conditions (and, for brevity, Table 7 summarizes the mean acquisition proportions accordingly).

TABLE 7 ABOUT HERE

Trials with one cue. The main effect of Cue was significant, $F(3,93) = 6.94$, $p < .001$, indicating that cues differed in how frequently they were examined. Tukey HSD post hoc tests revealed that cue C (greatest success/ECB and Pearson’s r) was examined significantly more frequently than each of the other three cues (all $p < .01$).

Trials with two cues. The main effect of Cue was significant, $F(3,93) = 3.63$, $p = .016$. Tukey post hoc tests revealed that Cue C was bought significantly more frequently than cue A (lowest Pearson's r and EV) ($p < .05$).

Trials with three cues. The main effect of Cue was significant, $F(3,93) = 3.50$, $p = .019$. Tukey post hoc tests revealed that Cue C was bought significantly more frequently than cue A ($p < .01$).

Proportions of Information Examined – the Effect of Cue and Amount of Information

Appropriate adaptation to the probabilistic structure of the environment would predict that participants would alter which cues they examine most when they are forced to examine fewer (or allowed to examine more) cues. Under such circumstances the Amount of information would interact with Cue, such that cues with high EV were more likely to be chosen (in comparison to other cues) when more cues could be examined. Two-way ANOVA was used to examine the pattern of cue examination as a function the Amount of information that participants could examine. The proportion of information examined was the dependent variable, and Cue (four levels, A to D) and Amount of information (three levels, one to three cues) were the independent variables. Amount of information is uninteresting in its own right (as proportions must increase if more cues are examined) – however including this factor permits examining the two-way interaction to determine whether patterns of cue examination change meaningfully as more information is examined.

Consistent with prior analyses, the main effect of Cue was significant, $F(3,93) = 5.89$, $p = .001$. Tukey HSD post hoc tests revealed that cue C (highest r and success/ECB) was examined more frequently than each of the other three cues ($p < .05$ for cues B and D; $p < .01$ for cue A). The two-way interaction was not significant, $F(6,186) = .972$, $p = .446$, thus

patterns of cue examination were relatively constant no matter how many cues could be examined.

Tallies of participants' most favored cue indicated that cue C (highest r and success/ECB) was the most frequently examined cue for 19.5 participants (61%). Cues A, B, and D were examined most frequently by 3.5, 5 and 4 participants respectively.

Patterns of cue examination were further examined using Page's nonparametric trend test (Table 7). For any number of cues acquired, cue examination is best described by Pearson's r , then success/ECB, then EV, and then DR. Acquisition proportions are significantly related to the orders prescribed by the Correlational and Bayesian frameworks, but not to other orders.

Thus, consistent with Experiments 1 and 2 there is clear evidence that the frequency with which a cue is examined is related to its 'importance' as defined by Correlational and Bayesian frameworks. However, there is no evidence favoring adaptation to the mix of DR and EV as a function of the certainty with which one can acquire information.

General Discussion

When making decisions we are often faced with the need to trade-off the costs – financial, cognitive or opportunity costs – of acquiring information against the perceived improvements in accuracy from obtaining extra information. We proposed that in multiple-cue probabilistic environments, adaptation in the acquisition of information could occur both at the 'gross' level of the amount of information acquired, and at the more subtle level of the 'informativeness' or value of acquired information. The results of three experiments suggest that people adapt appropriately at the 'gross' level, and, that although sensitive to the subtle variations in cue informativeness, do not necessarily alter patterns of acquisition in line with these variations. We consider these two aspects of adaptation in reverse order.

Adaptation to probabilistic structure: Which Framework is Best?

We took a novel and integrative approach by examining three separate frameworks for assessing the value of information carried by a cue: the Bayesian, Correlational, and Ecological approaches. Results indicated that each one has some explanatory power in accounting for the frequency with which cues are searched. Specifically, if cues are ordered from highest to lowest by the frequency with which they are examined, this ranking is related to the descending order of informativeness specified by each of these frameworks. All frameworks are of some value, but which is best? A reliable pattern was seen across the three experiments. The Correlational framework was best at describing patterns of search, followed by the Bayesian framework and then the Ecological framework (EV-directed search).

The different rank orders specified by the three frameworks were only subtly different. This was an inevitable feature of our design given that Pearson's r and success/ECB are monotonic increasing functions of EV and DR. Nonetheless, the ranking of explanatory power for these frameworks (Correlational > Bayesian > Ecological) was apparent in three experiments, using different cue-criterion environments and different task characteristics, suggesting that the results indicate a real and robust effect.

The implication of the success of the Correlational and Bayesian frameworks in predicting the search through cues implies that both Ecological Validity (EV) and Discrimination Rate (DR) are important features for determining how people search for information in probabilistic environments. The advantage of the Correlational framework over the Bayesian framework implies that the variation in EV across cues seems to carry greater impact than the variation in DR. (Pearson's r is more closely related to EV than is success/ECB. Success/ECB increases with EV and DR, whereas Pearson's r increases with EV and the *square root* of DR (Equations 1b and 2). Hence, variation in DR has less impact

upon Pearson's r than upon success/ECB, leaving the ranking of cues by Pearson's r more closely allied to that for EV.) EV can be thought of as the accuracy component of the information carried by the cue, whilst DR can be considered as the speed component (if cues with high DRs are examined, informed decisions can be made more quickly because such cues are more likely to discriminate between options). Thus the search through cues is sensitive to cue properties that enhance accuracy and reduce the amount of information that is considered. Our rational analyses show that if there is a cost to acquiring cues, both considerations are entirely appropriate.

The rational analyses for all experiments also indicated that a switch from an Ecological to a Correlational/Bayesian search order was advantageous as: a) the cost of information increased (Experiment 1), b) the time available for decisions decreased (Experiment 2) and c) the certainty about availability of information decreased (Experiment 3). Despite these prescribed changes in the optimal strategy we did not see concomitant shifts in the acquisition frequency measure. A plausible explanation for this resistance to shift can be framed in terms of setting an aspiration level for earnings. The manipulation of task characteristics (cue-cost, time, uncertainty) was within-subjects in the three experiments, all experiments involved numerous repeated choices, and participants' earnings were present on-screen throughout experimental sessions. The combination of these factors may have led participants to set a personal aspiration level and then to 'lock-in' to a particular pattern of acquisition that enabled them to maintain such a level. Indeed, one of the earliest choice heuristics proposed in the decision-making literature – 'satisficing' (Simon, 1955) - incorporates an aspiration level as a key component.

Because sub-optimal search orders only had a relatively small impact on over-all accuracy – and hence amount earned (e.g., see Figure 1), it is likely that a given aspiration level would have been unaffected by the switches between search orders. For example,

switching from the Ecological-Full to the Correlational-Full search strategy in the most expensive cue condition of Experiment 1 would have increased earnings by only 5.3 points per trial - a difference in magnitude of outcome that may only affect a few participants.

Much more likely to impact on a given aspiration level would be failure to adapt at the gross level of the raw amount of information acquired. For example buying all the cues instead of only the cue of highest EV on each trial (i.e., not even the optimal one in such circumstances) in the 24 point cue cost condition of Experiment 1 would lead to a decrease in earnings of 55.6 points per trial. Only the most ‘unambitious’ participant could afford to ignore a difference of this magnitude. Consistent with this explanation we did indeed find adaptation in the *amount* of information acquired under different cost conditions. We discuss below the theoretical implications of finding changes in the amount of information acquired, but first we consider the possibility that participants did not switch search strategies simply because they did not learn enough about individual cue parameters.

It is true that the differences between cues on our parameters of interest were often quite small. Search patterns were not identical across participants – which would most naturally be interpreted as evidence of considerable noise in cue learning resulting from the difficulty of distinguishing between cues. Nonetheless, participants’ patterns of search indicated that they were able to distinguish between cues. Group level tests of significance showed that participants distinguished between cues that differed by no more than 0.2 in their correlation with the criterion. In Experiments 2 and 3, cues that differed by only 0.1 in their correlation with the criterion were differentiated. We find this ability to differentiate between cues quite impressive, and consider it unsurprising that there should be some noise in learning which cues are more useful than others.

It remains to be seen whether any of the variation in search orders between different participants represents stable individual differences (as opposed to random noise in learning).

It is not inconceivable that some individuals will pay greater attention to the accuracy characteristics of cues (i.e., be more influenced by variation in EV), whilst others will be more concerned to enhance the speed of decisions (i.e., be more influenced by variation in DR). Consistent with such a view, in Experiments 1 and 2, the cue with highest EV but lowest DR showed the greatest variation in frequency of acquisition among the cues. This would be anticipated if participants have qualitatively different reactions to EV and DR. Such individual differences are plausible, as cognitive dispositions have been shown to explain significant variance in the performance of a range of thinking and reasoning tasks (Stanovich & West, 1998). Specifically, individual differences in impulsivity and willingness to seek disconfirming evidence described by Stanovich and West (1998) might account for some variation in the number of cues that participants examine. Such an effect could be mediated by a propensity to search for cues with higher, or lower, DRs.

Adaptation to Task Characteristics

In Experiments 1 and 2 we found evidence that the increase in the ‘cost’ of information led to a decrease in the average amount of information acquired. In Experiment 1 the cost was an explicit monetary cost and in Experiment 2 the ‘cost’ was associated with the time available to make a choice. (Recall that in Experiment 3 the amount of information available was pre-determined so looking at amount acquired *per se* was uninformative). Findings of an adaptation to cost structure are consistent with a wealth of literature on information acquisition under different task conditions (e.g., Bröder, 2000; Connolly & Serre, 1984; Newell & Shanks, 2003, Payne et al., 1993).

Despite the considerable amount of consistent evidence, there is tension between two distinct approaches to explaining the changes in the amount of information acquired when task conditions alter. On the one hand, changes in behavior are described in terms of

qualitative changes in the choice strategy or heuristic that is employed (Beach & Mitchell, 1978; Gigerenzer & Todd, 1999; Payne et al., 1993; Rieskamp & Hoffrage, 1999). Under such accounts there is a cost/benefit trade-off involved in selection between strategies that are more labour intensive (e.g. a weighted additive rule) and those that are more frugal but only sometimes less accurate (e.g., a one-reason choice strategy, Gigerenzer & Goldstein, 1996).

Alternatively, changes in behavior can be viewed as adaptations or adjustments in the application of a general or universal mechanism. A range of mechanisms has been proposed – many built up from models of simpler psychophysical discrimination tasks (e.g., Nosofsky & Palmeri, 1997; Ratcliff, 1978). While specific mechanisms differ, in general these models assume that rather than “taking” a predetermined quantity of information (e.g., one discriminating cue in the case of TTB), sequential sampling of each option occurs until evidence sufficient to favor one option over the other has been accumulated. For instance, Busemeyer and Townsend (1993) interpret the effect of increased time pressure on choice in terms of changes in the threshold that determines whether deliberation (information acquisition) continues. Lee and Cummins (2004) present a similar interpretation, finding that an evidence accumulation threshold model accounts for behavior better than TTB or a ‘rational’ model (which always examines every piece of information) in a two-alternative forced choice task.

Value and amount of information

The two types of adaptation we have examined can thus both be interpreted in terms of shifts between strategies or adjustments of thresholds. In the case of information content, the ‘switch’ can be conceptualized as occurring between search strategies that employ rules that differ in the relative influence of EV and DR. However, such changes could also be incorporated into a general model that searches for information according to a threshold of

information value. Similarly, as the preceding discussion shows, shifts in the *amount* of information acquired could reflect shifts between strategies that acquire more or fewer pieces of information, or a single mechanism, which accumulates evidence up to a threshold. Implicit in both interpretations is the notion of an aspiration level for total earnings in the experiment – an aspiration level that, in our experiments, appeared to be affected by the *amount* of information acquired but not by the *value* of that information.

Although with our data we are not able to discriminate conclusively between these alternative accounts, there are a number of theoretical reasons why one might favor a single adaptive model. A general mechanism with a threshold level provides a parsimonious account of why different amounts of information are examined under differing task conditions – it does not require the invention of a new strategy or heuristic to account for every new pattern of data. More importantly, by assuming that different people will require different amounts of information for the same decision it has greater potential for explaining individual variability under the same task conditions (Lee & Cummins, 2004). Besides parsimony a general model approach avoids the infinite regress problem, suffered by the multiple strategy approach, of needing a meta-heuristic to ‘choose’ the particular strategy or heuristic for a given task (Busemeyer, 1993; Newell, 2004). Finally, a general model can ‘mimic’ a variety of different strategies through the adjustment of an evidence threshold thus leaving little explanatory value to be added by a ‘toolbox’ of strategies (e.g., Bockenholt et al.; 1991 Einhorn, Kleinmuntz, & Kleinmuntz, 1979; Lee & Cummins, 2004).

A common rebuttal to the ‘single-mechanism’ approach is that the use of parameterised thresholds increases their complexity in comparison with simple-heuristic models, and it is this complexity, which leads to improvements in accuracy (e.g., Roberts & Pashler, 2000). In conclusion, we note that the evidence-accumulation model recently proposed by Lee and Cummins (2004) goes some way to formally defining the advantages of single models over

‘toolboxes’ of heuristics. Their model employs a Bayesian approach for estimating cue importance, combined with a random walk sequential sampling mechanism for accumulating evidence. Although the model has more parameters than a ‘one-reason’ strategy (e.g., TTB), it is more successful in capturing inter-participant individual differences as well as being preferred over TTB by a standard measure of model selection (minimum description length – see Pitt, Myung & Zhang, 2002). Given the success of the *Correlational* framework (over the Bayesian) in capturing our participants’ conceptions’ of the importance of cues it would be extremely useful to test a version of Lee and Cummins’ model with a correlational search rule under a variety of task conditions.

In this paper we have provided the means to integrate previous research on pre-decisional information search. We have shown that of the methods used previously for assessing the value of information, cue-outcome correlation best predicts *which cues* are acquired most frequently, though, perhaps inevitably, learning seems subject to noise. Rational analysis indicates that sensitivity to this characteristic is adaptive when information has some cost - as both the speed and accuracy with which decisions can be made is reflected by cue-outcome correlations. Participants were also sensitive to the need to alter the *number of cues* acquired in response to changes in task characteristics. This seems to have been achieved by truncating search, not by responding to more subtle features of the informativeness of cues that would have yielded yet better outcomes. Thus both the learning of cues-outcome relationships defined by the probabilistic structure of the environment, and the adaptation to task characteristics, were suboptimal – but may well be sufficient for the successful implementation of boundedly rational strategies.

References

- Albert, D., Aschenbrenner, K. M., & Schmalhofer, F. (1989). Cognitive choice processes and the attitude-behavior relation. In A. Upmeyer (Ed.), *Attitudes and Behavioral Decisions* (pp. 61-99). New York: Springer.
- Beach, L.R., & Mitchell, T.R. (1978). A contingency model for the selection of decision strategies. *Academy of Management Review*, 3, 439-449.
- Bockenholt, U., Albert, D., Aschenbrenner, M., & Schmalhofer, F. (1991). The effects of attractiveness, dominance and attribute differences on information acquisition in multiattribute binary choice. *Organizational Behavior and Human Decision Processes*, 49, 258-281.
- 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.
- Brunswik, E. (1955). Representative design and probabilistic theory in a functional psychology. *Psychological Review*, 62, 193-217.
- Busemeyer, J. (1993). Violations of the speed accuracy trade-off relation: Decrease in decision accuracy with increases in time. In O. Svenson & J. Maule (Eds.) *Time pressure and stress in human judgment and decision making*. (pp 181-193). New York: Plenum Press.
- Busemeyer, J. R., & Townsend, J. T. (1993). Decision Field Theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100, 432-459.
- Cohen, J. (1969). *Statistical Power Analysis for the Behavioral Sciences*. New York: Academic Press.

- Connolly, T., & Serre, P. (1984). Information search in judgment tasks: The effects of unequal cue validity and cost. *Organizational Behavior and Human Performance*, 34, 387-401.
- Edwards, W. (1965). Optimal strategies for seeking information. *Journal of Mathematical Psychology*, 2, 312-329.
- Einhorn, H. J., Kleinmuntz, D. N., & Kleinmuntz, B. (1979). Linear regression and process-tracing models of judgement. *Psychological Review*, 86, 465-485.
- Fried, L. S., & Peterson, C. R. (1969). Information seeking: optional versus fixed stopping. *Journal of Experimental Psychology*, 80, 525-529.
- Gigerenzer, G., & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological Review*, 103, 650-669.
- Gigerenzer, G., & Todd, P. M. (Eds.). (1999). *Simple Heuristics That Make Us Smart*. New York: Oxford University Press.
- Hammond, K. R., & Stewart, T. R. (Eds.). (2001). *The Essential Brunswik: Beginnings, Explications, Applications*. New York: Oxford University Press.
- Holzworth, J. R. (2001). Multiple cue probability learning. In K. R. Hammond & T. R. Stewart (Eds.), *The Essential Brunswik: Beginnings, Explications, Applications*. (pp. 348-350). New York: Oxford University Press.
- Howell, D. C. (2002). *Statistical Methods for Psychology (5th edition)*. Pacific Grove, CA: Duxbury Press.
- Klayman, J., & Ha, Y.-W. (1987). Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, 94, 211-228.
- Lee M.D. and Cummins T.D.R. (2004). Evidence accumulation in decision making: Unifying the 'Take-the-best' and the 'rational' models. *Psychonomic Bulletin & Review*, 11, 343-352.

- Lindley, D. V. (1956). On a measure of the information provided by an experiment. *Annals of Mathematical Statistics*, 27, 986-1005.
- Martignon, L., & Hoffrage, U. (1999). Why does one-reason decision making work? A case study in ecological rationality. In G. Gigerenzer, & P. M. Todd (Eds.), *Simple Heuristics That Make Us Smart* (pp. 119-140). New York: Oxford University Press.
- Newell, B.R. (2004). Re-visions of rationality? Manuscript submitted for publication, University of New South Wales, Sydney, Australia.
- Newell, B. R., Rakow, T., Weston, N. J., & Shanks, D. R. (2004). Search strategies in decision making: The success of success. *Journal of Behavioral Decision Making*, 17, 117-137.
- 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.
- Newell, B. R., Weston, N. J., & Shanks, D. R. (2003). Empirical tests of a fast and frugal heuristic: Not everyone "takes-the-best". *Organizational Behavior and Human Decision Processes*, 91, 82-96.
- Nosofsky, R.M. and Palmeri, T.J. (1997). An exemplar-based random walk model of speeded classification. *Psychological Review*, 104, 266-300.
- Oaksford, M., & Chater, N. (1998). *Rationality in an Uncertain World: Essays on the Cognitive Science of Human Reasoning*. Hove, UK: Psychology Press.
- Payne, J. W. (1982). Contingent decision behavior. *Psychological Bulletin*, 92, 382-402.
- Payne, J. W., Bettman, J. R., & Johnson, E. J. (1993). *The Adaptive Decision Maker*. New York: Cambridge University Press.
- Pitt, M.A., Myung, I.J, & Zhang, S. (2002). Toward a method of selecting among computational models of cognition. *Psychological Review*, 109, 472-491.

- Rakow, T., Hinvest, N., Jackson, E., & Palmer, M. (2004). Simple Heuristics from the Adaptive Toolbox: Can we perform the requisite learning? *Thinking & Reasoning*, 10, 1-29.
- Rapoport, A., & Wallsten, T. S. (1972). Individual decision behavior. *Annual Review of Psychology*, 23, 131-179.
- Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85, 59-108.
- Rieskamp, J., & Hoffrage, U. (1999). When do people use simple heuristics, and how can we tell? In G. Gigerenzer & P. Todd (Eds.), *Simple Heuristics That Make Us Smart* (pp. 141-167). New York: Oxford University Press.
- Roberts S., & Pashler, H. (2000). How persuasive is a good fit? A comment on theory testing. *Psychological Review*, 107, 358-367.
- Siegel, S., & Castellan, N. J. (1988). *Nonparametric statistics for the behavioral sciences (Second Edition)*. New York: McGraw-Hill.
- Simon, H.A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 99-118.
- Slovic, P., & Lichtstein, S. (1971). Comparison of Bayesian and regression approaches to the study of information processing in judgement. *Organizational Behavior and Human Performance*, 6, 649-744.
- Smedslund, J. (1955). *Multiple-Probability Learning*. Oslo: Akademisk Forlag.
- Stanovich, K. E., & West, R. F. (1998). Individual differences in rational thought. *Journal of Experimental Psychology: General*, 127, 161-188.
- Todd, P. M., & Gigerenzer, G. (2000). Précis of Simple Heuristics That Make Us Smart. *Behavioral & Brain Sciences*, 23, 727-780.
- Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, 79, 281-299.

Author Note

The first and second authors contributed equally to this project and order of authorship was determined arbitrarily. We gratefully acknowledge the assistance of Dr Roger Grace who programmed Experiments 2 and 3, and the financial support of the Economic and Social Research Council (ESRC) and the Leverhulme Trust. The work was part of the programme of the ESRC Research Centre for Economic Learning and Social Evolution. Correspondence concerning this article should be addressed to Tim Rakow, Department of Psychology, University of Essex, Wivenhoe Park, Colchester, Essex, CO4 3SQ (Email: timrakow@essex.ac.uk) or Ben Newell, School of Psychology, University of New South Wales, Sydney, 2052, Australia (Email: ben.newell@unsw.edu.au).

Table 1. Multiple-cue probability-learning choice task: Assessing insurance risk from imperfect predictors.

<i>Cue</i>	<i>Is the client...</i>	Client A	Client B
<i>Sex</i>	<i>Male?</i>	Yes	No
<i>Age</i>	<i>Under 25 years of age?</i>	??	??
<i>Previous convictions</i>	<i>A convicted felon?</i>	No	No
<i>Marital status</i>	<i>Married?</i>	??	??

Table 2. The experimental environment in Experiment 1

	Cue A	Cue B	Cue C	Cue D	Defines the search order	Strategy searching in this order
DR	.219	.375	.500	.469	C > D > B > A	--
EV	.825	.606	.671	.768	A > D > C > B	Ecological
Success/ECB	.571	.540	.586	.626	D > C > A > B	Bayesian
Pearson's <i>r</i>	.304	.130	.242	.367	D > A > C > B	Correlational

DR = discrimination rate, EV = ecological validity, ECB = expected change in belief

Table 3. Experiment 1: (a) Mean (SD) proportion of trials on which each cue is purchased, and (b) z for specified trends in cue proportions purchased

Cost condition	(a) Proportions purchased						(b) Order examined by Page Trend Test			
	Cue A	Cue B	Cue C	Cue D	All cues	Search frequency	DR	EV	Success/ECB	Pearson r
3 points	.65 (.43)	.32 (.29)	.56 (.29)	.83 (.20)	.59 (.18)	D > A > C > B	0.50	2.50**	2.70**	3.00**
10 points	.49 (.40)	.18 (.19)	.47 (.27)	.83 (.28)	.49 (.13)	D > A > C > B	0.85	2.45**	3.15***	3.30***
17 points	.52 (.45)	.09 (.12)	.42 (.34)	.61 (.34)	.41 (.13)	D > A > C > B	0.45	2.95**	2.75**	3.10***
24 points	.46 (.48)	.11 (.26)	.33 (.37)	.51 (.47)	.35 (.11)	D > A > C > B	0.50	1.70*	1.90*	2.00*
Combined	.53 (.40)	.18 (.17)	.44 (.24)	.69 (.25)		D > A > C > B	0.90	2.30*	3.10***	3.20***

DR = discrimination rate, EV = ecological validity, ECB = expected change in belief

* $p < .05$, ** $p < .01$, *** $p < .001$

Table 4. The experimental environment in Experiment 2

	Cue A	Cue B	Cue C	Cue D	Defines the search order	Strategy searching in this order
DR	.469	.375	.500	.219	C > A > B > D	--
EV	.776	.608	.677	.840	D > A > C > B	Ecological
Success/ECB	.629	.541	.589	.574	A > C > D > B	Bayesian
Pearson's <i>r</i>	.377	.134	.252	.316	A > D > C > B	Correlational

DR = discrimination rate, EV = ecological validity, ECB = expected change in belief

Table 5. Experiment 2: (a) Mean (SD) proportion of trials on which each cue is examined, and (b) z for specified trends in cue proportions examined

Time pressure condition	(a) Proportions examined						(b) Order examined by Page Trend Test			
	Cue A	Cue B	Cue C	Cue D	All cues	Search frequency	DR	EV	Success/ECB	Pearson r
Low	.87 (.22)	.44 (.25)	.57 (.27)	.61 (.37)	.55 (.08)	A > D > C > B	1.13	2.97**	4.50***	4.78***
High	.81 (.26)	.36 (.29)	.54 (.25)	.56 (.34)	.62 (.13)	A > D > C > B	1.29	2.84**	4.23***	4.32***
Combined	.84 (.22)	.40 (.24)	.56 (.25)	.59 (.35)		A > D > C > B	1.10	2.89**	4.29***	4.53***

DR = discrimination rate, EV = ecological validity, ECB = expected change in belief

* $p < .05$, ** $p < .01$, *** $p < .001$

[z -values differing by more than 1.96 are significantly different ($p < .05$) from each other, hence purchase proportions are significantly more closely related to success/ECB or Pearson's r than to DR]

Table 6. The experimental environment for Experiment 3

	Cue A	Cue B	Cue C	Cue D	Defines the search order	Strategy searching in this order
DR	.494	.198	.444	.346	A > C > D > B	--
EV	.697	.829	.785	.747	B > C > D > A	Ecological
Success/ECB	.597	.565	.627	.585	C > A > D > B	Bayesian
Pearson's <i>r</i>	.276	.292	.381	.289	C > B > D > A	Correlational

DR = discrimination rate, EV = ecological validity, ECB = expected change in belief

Table 7. Experiment 3: (a) Mean (SD) proportion of trials on which each cue is examined, and (b) z for specified trends in cue proportions examined

Amount of information	(a) Proportions purchased					(b) Order examined by Page Trend Test			
	Cue A	Cue B	Cue C	Cue D	Search frequency	DR	EV	Success/ECB	Pearson r
One cue	.16 (.23)	.19 (.29)	.48 (.38)	.17 (.24)	C > B > D > A	0.74	0.92	2.51**	2.63**
Two cue	.38 (.35)	.43 (.34)	.68 (.33)	.47 (.37)	C > D > B > A	0.12	1.23	1.90**	2.63**
Three cue	.59 (.34)	.70 (.29)	.86 (.22)	.70 (.30)	C > B > D > A	0.37	1.38	2.42**	3.09**
All trials	.38 (.27)	.44 (.26)	.68 (.28)	.45 (.25)	C > D > B > A	0.34	1.53	2.60**	3.40***

DR = discrimination rate, EV = ecological validity, ECB = expected change in belief

* $p < .05$, ** $p < .01$, *** $p < .001$

Figure 1. Expected payoff per trial for different strategies (100 points for a correct choice)

[The uppermost line corresponding to any point along the horizontal axis indicates the best strategy for a given cue cost.]

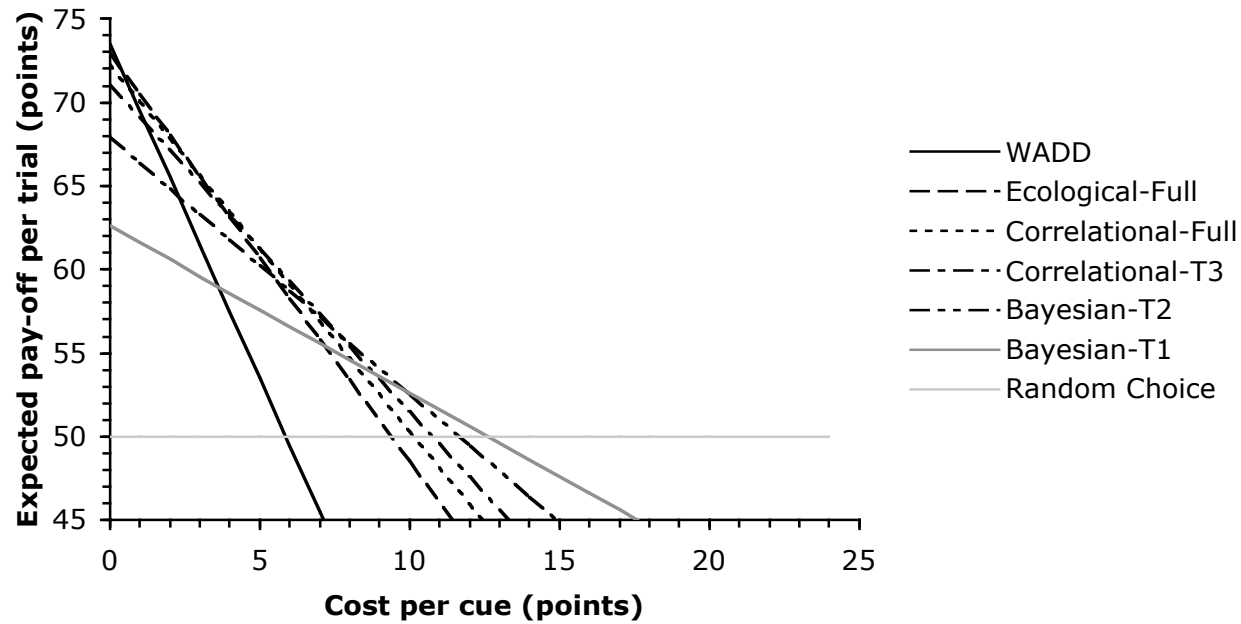
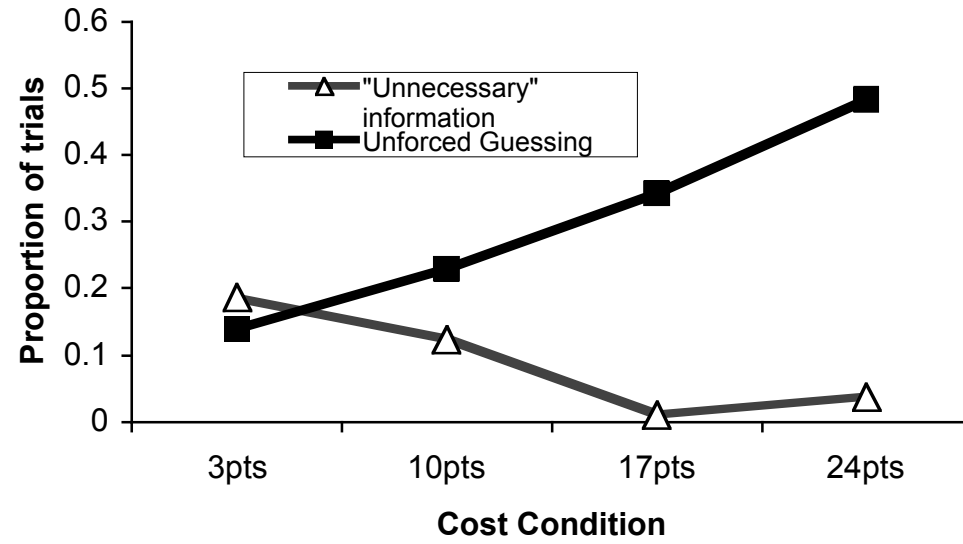


Figure 2. Departures from a one-reason stopping rule with varying cue cost.



Appendix 1. Expressing Pearson's r as a function of ecological validity (EV) and discrimination rate (DR)

Definitions. For a cue with discrimination rate (d) and ecological validity (v), the success (s) can be given by: $s = (v - 0.5)d + 0.5$

Purpose. Below we demonstrate that Pearson's r and success define similar hierarchies of cues, as it can be shown that:

$$r = \frac{2s - 1}{d}$$

or equivalently that

$$r = (2v - 1)(d)$$

To demonstrate this, we first derive an intermediate result (Result A below).

Demonstration of Result A. Consider P paired choices between objects A and B, where R correct choices and W incorrect choices are made. If objects A and B are defined arbitrarily, we expect both correct and incorrect choices to be divided equally between the two objects. Define X as the direction the cue points to, which takes values $+1$ (cue points to object A), 0 (cue points to neither object), or -1 (cue points to object B). Define Y as the outcome of the target event, which takes values $+1$ (object A has the higher criterion value) or -1 (object B has the higher criterion value).

The general expression for Pearson's r is (Siegel & Castellan, 1988):

$$r = \frac{\sum xy}{(\sum x^2 \sum y^2)^{1/2}} \tag{1}$$

where, $x = X - \bar{X}$, and $y = Y - \bar{Y}$

By symmetry both \bar{X} & \bar{Y} are 0, giving the following cross-tabulated frequencies for x and y:

Values of x	Values of y		
	+1	-1	Total
+1	R/2	W/2	(R+W)/2
0	(P-R-W)/2	(P-R-W)/2	(P-R-W)
-1	W/2	R/2	(W+R)/2
Total	P/2	P/2	P

Using (1):

$$r = \frac{(R/2)(+1) + ((P-R-W)/2)(0) + (W/2)(-1) + (W/2)(-1) + ((P-R-W)/2)(0) + (R/2)(+1)}{**\{(P/2)(+1)^2 + (P/2)(-1)^2\}}$$

$$r = \frac{(R/2) - (W/2) - (W/2) + (R/2)}{\{(R+W)/2 + (W+R)/2\}\{(P/2) + (P/2)\}}$$

$$r = \frac{R - W}{(P(R + W))} \quad \text{Result A}$$

Result A also defines Spearman's rank correlation coefficient (ρ , or *rho*) for the environments that we consider, as ρ can be calculated by applying the formula for Pearson's *r* to the ranks of the scores (Howell, 2002; Siegel & Castellan, 1988). (*r* and ρ would only differ were we to score the direction that the cue points to in some asymmetrical fashion, presumably reflecting that pointing to object A were somehow more, or less, significant than pointing to object B. We can think of no instance where this would seem reasonable.)

Demonstration of main result. Starting with Martignon and Hoffrage's (1999)

definition of success in terms of the number of correct (R) and incorrect (W) choices from a set of P pairs:

$$s = \frac{R + 0.5(P - R - W)}{P}$$

$$\square \quad 2s = \frac{2R + P - R - W}{P}$$

$$\square \quad 2s = \frac{P + R - W}{P}$$

$$\square \quad 2s = 1 + \frac{R - W}{P}$$

$$\square \quad P(2s - 1) = R - W \quad (2)$$

Substituting (2) into Result A:

$$r = \frac{P(2s - 1)}{(P(R + W))}$$

$$\square \quad r = \frac{(2s - 1) P}{(R + W)} \quad (3)$$

By definition, the discrimination rate is:

$$d = \frac{R + W}{P}$$

$$\square \quad \frac{1}{d} = \frac{P}{(R + W)} \quad (4)$$

Substituting (4) into (3) above:

$$r = \frac{2s - 1}{d} \quad (5a)$$

Using the definition of success above: $s = (v - 0.5)d + 0.5$, and substituting this into (5a):

$$r = \frac{2\{(v - 0.5)d + 0.5\} - 1}{d}$$

$$r = \frac{(2v - 1)d + 1 - 1}{d}$$

$$r = (2v - 1)(d) \quad (5b) \text{ [equation (2) in the Introduction]}$$

Footnotes

1. There were no significant effects of cost condition order or session order on the proportion of information purchased.
2. As payment in this experiment was determined by the probability of outcomes, and not the outcome itself (which is subject to random variation), no feedback was given on the outcome of each trial. Feedback of the probability that the chosen share was the better share was retained.
3. Though statistically significant, the mean difference between conditions in the proportion of the available information examined (equivalent to 0.27 cues per trial) was smaller than we had anticipated. Notably, the mean (SD) time taken per trial was 7.61 (1.62) seconds under Low time pressure, and 9.05 (1.21) under High time pressure. Given that the time limit given to participants allowed them to take an average of 10 seconds per trial, there was a clear tendency for participants to collect less information than the time permitted them to in the Low time pressure condition.
4. With respect to the manipulation of cue cost, this manipulation is a little different to those employed in the first two experiments. In Experiments 1 and 2, cue costs were manipulated (financially, or via time pressure) and we *observed* the effect of this upon cue acquisition, and found that levels of cue acquisition varied considerably between individuals. One way to think of this limitation to one, two, or three cues is by specifying that it is *as if* cue costs have been manipulated such that participants opt to select one, two, or three cues on a given trial. (The costs required to achieve this would inevitably vary from one participant to another.)