

Multi-attribute utility models as cognitive search engines

Pantelis P. Analytis*

Amit Kothiyal†

Konstantinos Katsikopoulos†

Abstract

In optimal stopping problems, decision makers are assumed to search randomly to learn the utility of alternatives; in contrast, in one-shot multi-attribute utility optimization, decision makers are assumed to have perfect knowledge of utilities. We point out that these two contexts represent the boundaries of a continuum, of which the middle remains uncharted: How should people search intelligently when they possess imperfect information about the alternatives? We assume that decision makers first estimate the utility of each available alternative and then search the alternatives in order of their estimated utility until expected benefits are outweighed by search costs. We considered three well-known models for estimating utility: (i) a linear multi-attribute model, (ii) equal weighting of attributes, and (iii) a single-attribute heuristic. We used 12 real-world decision problems, ranging from consumer choice to industrial experimentation, to measure the performance of the three models. The full model (i) performed best on average but its simplifications (ii and iii) also had regions of superior performance. We explain the results by analyzing the impact of the models' utility order and estimation error.

Keywords: ordered search, optimal stopping, linear models, subjective utility, heuristics

1 Introduction

Do the following thought experiment: You are the human resources manager of a company and you are assigned the task of hiring a new employee. After advertising the position, you receive several dozen applications from candidates listing their skills and credentials (e.g., grade point average, work experience, programming skills). You can determine each candidate's potential only after inviting him or her for an interview. Let us assume that you can interview candidates sequentially and that you can decide to stop interviewing and hire a candidate after each interview. Crucially, making the effort to interview another candidate is costly. What is the best way to organize the interview process? First, you need to decide the order in which you will be inviting candidates. Then, after each interview you need to decide whether to make an offer to one of the interviewed candidates, thus stopping your search. The first problem is an ordering problem and the second a stopping problem.

Clearly, if you could perfectly estimate the potential of the candidates on the basis of their credentials you could directly choose the best one by using decision an-

alytic methods (e.g., Keeney & Raiffa, 1993). On the other hand, if the credentials were not at all informative, you would have to invite people at random, and your problem would reduce to an optimal stopping problem. Such models have been developed formally in statistics (DeGroot, 1970) and economics (Stigler, 1961; Lippman & McCall, 1976) and human behavior in them has been tested empirically in psychology (Rapoport & Tversky, 1966; Lee, 2006), economics (Schotter & Braunstein, 1981; Hey, 1987; Sonnemans, 1998) and marketing (Zwick, Rapoport, Lo & Muthukrishnan, 2003). Intuitively, most everyday decision-making problems lie between these two extreme cases; in reality, the attributes of the alternatives can be used to predict their utility but only imperfectly. There often remains some amount of uncertainty that cannot be explained by the attributes.

There has been some work on ordered search: In a seminal paper, Weitzman (1979) put forward a general model for ordering alternatives and terminating search in search problems with recall, where decision makers initially have partial information about the alternatives but learn their true utility after paying a search cost. Although Weitzman's results readily generalize to scenarios with multi-attribute alternatives, only a recent study by Dzyabura (2013) considered explicitly how sequential search can be guided by a multi-attribute utility model. Earlier, Roberts and Lattin (1991) presented a model of consideration set formation in which the decision makers include alternatives in their consideration set guided by a compensatory multi-attribute utility model. In essence, Roberts and Lattin's model can be seen as an ordered search model, in

We would like to thank Alexia Delfino, Juliane Kämmer, Hrvoje Stojic, Peter Todd and Jack Willis for their remarks on earlier versions of the paper and the members of the ABC group for their feedback during internal presentations. We would also like to thank Anita Todd for editing the manuscript.

Copyright: © 2014. The authors license this article under the terms of the Creative Commons Attribution 3.0 License.

*Max Planck Institute for Human Development, Lentzeallee 94, 14195, Berlin, Germany. E-mail: analytis@mpib-berlin.mpg.de.

†Max Planck Institute for Human Development.

which the number of alternatives that will be searched has to be decided once and for all before any search is performed.¹ However, the authors did not connect their results to search theory. Further, Moorthy, Ratchford and Talukdar (1997) employed Weitzman's model to derive predictions about the length of search in consumer choice. They further specified the original model and assumed that decision makers are uncertain about a brand's utility. The consumer beliefs in their model are probabilistic, and more experienced consumers are better at differentiating brands in terms of utility. Moorthy et al. tested their predictions on data from the car market. Last, in economics, several papers started from the assumption that agents search the alternatives in an externally imposed (Arbatskaya, 2007) or subjectively defined (Bagwell & Ramey, 1994; Armstrong, Vickers & Zhou, 2009) order and studied the aggregate market behavior.

In this paper we show that, when the decision makers' preferences can be described by a linear utility model, the problems introduced by Weitzman have an intuitive and psychologically plausible solution. Returning to our example, we will analytically show that the optimal policy is to follow the estimated utility order prescribed by your subjective utility model; then stop when the expected return from seeing one more candidate for the job turns negative. In essence, the utility models play the role of cognitive search engines, generating the order in which alternatives are examined. We formally develop this approach and apply it to three models that have been studied extensively in the field of judgment and decision making: (i) multi-attribute linear utility, (ii) equal weighting of attributes and (iii) a single-attribute heuristic. The simpler models (ii) and (iii) have been shown to perform well under some conditions (e.g., scarce information available for calibrating models) in one-shot choice problems (Baron, 1987; Gigerenzer, Todd & the ABC Research Group, 1999; Fasolo, McClelland & Todd, 2007; Katsikopoulos, 2011). We then compare the performance of the models in 12 real-world environments ranging from consumer choice to industrial experimentation and examine how the models' expected utility order and estimation error influence their performance and length of search.

Conceptually, our approach illustrates that optimal stopping problems assuming random search and one-shot choice problems are the boundary cases of an ordered search problem with imperfect information. It provides a formal framework within which the assumptions made by ordered search models, such as those proposed by Bagwell and Ramey (1994), Moorthy et al. (1997) and Arm-

¹This was the case in the first search models that were introduced in economics (e.g. Stigler, 1961). Those models were later called fixed-sample-size models. Morgan and Manning (1985) discussed in detail the points of divergence between fixed-sample-size models and sequential sampling models, in which the decision maker can stop search at any search step.

strong, Vickers and Zhou (2009), can be further clarified. In practice, our approach extends discrete choice models by specifying the exact search process. It is a plausible alternative to Roberts and Lattin's (1991) theory of consideration set formation and it further advances our understanding of decision-making in environments with rank-ordered alternatives.

In what follows, as in the model presented by Weitzman (1979), we focus on a scenario in which decision makers learn the exact utility of an alternative after sampling it and can always choose alternatives that they have sampled in the past. In Section 2 we develop a formal framework of optimal ordering and stopping. In Section 3 we test three models in 12 real-world environments ranging from consumer choice to industrial experimentation and examine the ordering and estimation error components of the models. In Section 4 we discuss the possibility of connecting our findings to other work, assess the conceptual and applied implications of our approach, and finally discuss the limitations and possible extensions of our framework.

2 The theoretical framework

2.1 The environment

There are n alternatives A_1, \dots, A_n . Each alternative A_i , $i \in \{1, \dots, n\}$ is associated with a vector of attributes $\mathbf{a}_i = (a_{i1}, \dots, a_{ik})$ and a utility u_i of choosing it. The u_i s are unknown but the \mathbf{a}_i s are known to the decision maker. The decision maker estimates u_i by $f(\mathbf{a}_i)$. We assume that the estimation errors, ϵ_i , such that $u_i = f(\mathbf{a}_i) + \epsilon_i$, are iid Gaussian with mean μ and standard deviation σ . We call this equation the decision maker's subjective model. For each alternative, the decision maker can only learn the utility u_i by sampling the alternative and paying a cost c . The decision maker can sample as many alternatives as desired, and she can eventually choose only among the sampled ones. If the decision maker searches for k items, the cost to be paid is kc and the decision maker will choose out of these k the alternative with the highest utility.

2.2 The optimal strategy

Let S denote the set of alternatives already searched and \bar{S} denote the set of alternatives not searched yet. That is,

$$S \cup \bar{S} = \{A_1, \dots, A_n\}, \quad S \cap \bar{S} = \emptyset$$

The decision maker's problem is to determine the search order and the stopping rule. Let the variable y denote the maximum utility that the decision maker can obtain from the alternatives in S , $y = \max_{A_k \in S} u_k$, where A_k belongs to S . If the decision maker sampled just one more item A_k before stopping search then the subjective expected gain

(i.e. the increase in utility minus cost) is (probabilities and expectations below are based on the decision maker's subjective model):

$$\begin{aligned}
 R(A_k) &= P(u_k > y) \times E(u_k - y - c | u_k > y) \\
 &\quad - P(u_k \leq y) \times c \\
 &= P(u_k > y) \times E(u_k - y | u_k > y) - c. \tag{1}
 \end{aligned}$$

It is intuitive that the decision maker should keep on sampling as long as there exists an alternative $A_k \in \bar{S}$ such that its subjective expected gain $R(A_k) > 0$. (If all $R(A_k) < 0$, search should be stopped.) Given this, the decision maker should sample the alternative A_k that achieves the maximum subjective gain $R(A_k)$. It turns out that to maximize $R(A_k)$, it suffices to select the alternative with the highest $f(\mathbf{a}_i)$ (for a proof, see Appendix 1):

Result 1. *If for two alternatives $A_i, A_j \in \bar{S}$, $f(\mathbf{a}_i) > f(\mathbf{a}_j)$ then $R(A_i) > R(A_j)$.*

This result says that if the decision maker decides to sample one more item before terminating the search, then the choice should be the one that has maximum unconditional expectation $E(u_i) = f(\mathbf{a}_i)$. This suggests the following policy:

Selection rule: Order the alternatives based on their unconditional expectation. Select the items for sampling in this order.

Stopping rule: If at any stage subjective expected gain is negative, terminate the search.

Note that the stopping rule can be applied only if the standard deviation of the estimation error σ is estimated. On the other hand, for the selection rule to be applied, only the parameters of the multi-attribute function $f(\mathbf{a}_i)$ need to be estimated. For $\sigma = 0$ the decision-making problem reduces to a single choice. For $c = 0$ the decision maker searches all the alternatives. Note that for $c = 0$ and when there are only two alternatives the model reduces to the probit model.² If $\sigma = \infty$, $P(u_i > y)$ is equal to 0.5.

2.3 Subjective models

In our framework, the order of the alternatives A_i is based on their subjective utility $f(\mathbf{a}_i)$. Thus, the decision maker's eventual success depends upon the multi-attribute utility function he or she uses to estimate utilities and generate an order of the alternatives. We present three psychologically plausible and widely used multi-attribute utility functions:

²Then, for two alternatives A_i and A_j the probability that the alternative A_i will be chosen can be written as $\Phi\left(\frac{f(\mathbf{a}_i) - f(\mathbf{a}_j)}{\sqrt{(2)\sigma}}\right)$.

1. Multi-attribute linear utility (MLU):

$$f((a_{i_1}, \dots, a_{i_m})) = \sum_j \beta_j a_{i_j}.$$

MLU is one of the cornerstone models in research on multi-attribute decision making (Keeney & Raiffa, 1993). It is also the model used to derive consumer preferences in conjoint analysis surveys in marketing. MLU is the equivalent of multiple linear regression, which has been widely studied as a model of inductive inference in multi-cue learning (Hammond, Hursch & Todd, 1964).

2. Equal-weighted linear utility (EW):

$$f((a_{i_1}, \dots, a_{i_m})) = \sum_j a_{i_j},$$

where all the attributes a_{i_1}, \dots, a_{i_m} are normalized and brought to the same scale.

EW is a special case of MLU where all decision weights β_j s are equal. It was originally proposed as an alternative to multi-linear regression by Dawes and Corrigan (1974).

3. Single-attribute utility (SA):

$$f((a_{i_1}, \dots, a_{i_m})) = a_{i_j},$$

where a_{i_j} has the highest ecological validity among $\{a_{i_1}, \dots, a_{i_m}\}$ as expressed by Kendall's tau non-parametric correlation.


Versions of the SA model have been studied by Hogarth and Karelaia (2005). The SA model is akin to the lexicographic heuristic (Payne, Bettman & Johnson, 1993; Kohli & Jedidi, 2007) and the take-the-best heuristic (Gigerenzer & Goldstein, 1996; Katsikopoulos, Schooler & Hertwig, 2010). However, SA resolves ties between alternatives by choosing at random, whereas the lexicographic heuristic and take-the-best examine additional attributes.

Note that the three presented models can also be seen as cardinal estimation models, where the criterion value corresponds to the utility. Davis-Stober, Dana and Budescu (2010b) have pointed out that the notation of utility and estimation models can be used interchangeably.

2.4 An example

We illustrate how the different components of equation 1 play out with a concrete example (see Figure 1). Consider a scenario where a decision maker is searching in an online store to buy a single album of an up-and-coming music band she just heard about on the radio. The band has produced three albums so far. The decision maker can learn the exact utility of an album by listening to its songs. Her subjective beliefs are described by the SA model. As represented by the straight line in Figure 1, the decision maker believes that the expected utility of an alternative can be estimated by $u_i = 0.3 * a_i$, where a_i is the average rating of the album by other users of the site. As represented by the bell-shaped curves, she believes that the estimation error ϵ_i of her model is iid Gaussian with mean $\mu = 0$ and standard deviation $\sigma = 0.5$. The decision maker first samples album K, which has the high-

Figure 1: An example of the decision-making process. There are three alternatives available in the market represented by the points *K*, *L* and *M* on the plot. The average rating a_i (x axis) maps to the subjective utility (y axis) of the alternatives for the decision maker but with some error. As represented by the straight line, the decision maker believes that the expected utility of the alternatives can be estimated by $u_i = 0.3 * a_i$ where a_i is the average rating. As represented by the bell-shaped curves, the decision maker believes that the estimation error ϵ_i of her model is iid Gaussian with mean $\mu = 0$ and standard deviation $\sigma = 0.5$. The inset table shows the expected utility (EU) and realized utility (RU) of the alternatives. The decision maker will first sample the alternative with the highest attribute value and then at each step decide whether to sample the next alternative. For $c = 0.05$, for example, the decision maker will sample the alternatives *K* and *L* and choose the alternative *L*.



est expected utility. She finds out that the utility of album *K* is 1.87, which is slightly less than its expected utility. Then she has to decide whether it is worthwhile to examine the album *L*. Following equation 1 the expected returns from sampling album *L* can be written as $P(u_L > u_K) \times E(u_L - u_K | u_L > u_K) - c$. $P(u_L > u_K) = 0.33$, $E(u_L - u_K | u_L > u_K) = 0.329$ and their the product equals 0.109.³ Thus, the decision maker will examine the second album if the cost of search is lower than 0.109; otherwise she will stop search, choose album *K*, and never learn the actual utility of album *L*. Let us

³Note that $\sigma = 0.5$ appears in both of these two terms. If the decision maker believed $\sigma = 1$ the $P(u_L > u_K)$ would increase to 0.413 and $E(u_L - u_K | u_L > u_K)$ to 0.723 and the total benefits would be 0.299, almost three times higher than in the case where $\sigma = 0.5$.

assume that the cost is 0.05. Then, the decision maker samples album *L*. She finds out that the utility of *L* for her is 2.27, which is higher than her expectation and the utility of album *K*. Thus, *L* replaces *K* as the sampled album with the highest utility (y in equation 1). Now the returns from sampling album *M* can be written $P(u_M > u_L) \times E(u_M - u_L | u_M > u_L) - c$. $P(u_M > u_L) = 0.003$ and $E(u_M - u_L | u_M > u_L) = 0.152$. The product of these two parts equals 0.0005. Thus the overall return is negative and the decision maker will stop search after sampling album *L* and choose it. She will never learn the realized utility of album *M*.

3 Results

We applied the three models to study guided search in real-world problems. We examined the performance of the models in 12 data sets ranging from consumer choice to industrial experimentation. All the datasets included a variable with positive and more-is-better valence, which we treated as the utility. All the attribute values were normalized to a 0–1 scale to implement the EW model. When a correlation between an attribute and the criterion value in the test data set was positive, we converted the lowest attribute value to 0 and the highest to 1. When a correlation was found to be negative we converted the highest attribute value to 0 and the lowest to 1. We used ordinary least squares to calibrate the parameters β_j of the three models. For each of the models, we estimated the standard deviation of the error component σ using the standard error of the corresponding linear regression model $\hat{\sigma} = \sqrt{\frac{\sum_{t=0}^T (y - \hat{y})^2}{T}}$ where T is the number of alternatives in the training set and \hat{y} the estimate of the linear regression model with parameters β_j .

3.1 The environments

The 12 datasets analyzed are freely available from online databanks. We provide an additional reference list of the publications in which the datasets were originally reported in the supplementary material. As reported in Table 1, the datasets are characterized by a diverse number of attributes, alternatives, and intercorrelations between the attributes. We normalized the utility variable, setting the utility of the alternative with the lowest value equal to 0 and that with the highest value equal to 1. This transformation was necessary in order to achieve comparability across the datasets. The remaining variables are the attributes that were used to predict the alternatives' utility. In a few cases we excluded variables that were not suitable as attributes in a choice problem. In the white wine quality and the red wine quality environments we reduced the number of alternatives to 200 by choosing once at ran-

Table 1: The environment names refer to the variable that was assumed to be the utility in that environment. Asterisks indicate that the source data set was composed of several alternatives that were available in the market at the point of data collection. All the remaining data sets represent the results of controlled experiments. No. A refers to the number of alternatives and No. a to the number of attributes. The notation $|max(\rho_{ua_i})|$ indicates the strongest Pearson correlation between the utility variable and the attributes, $|\bar{\rho}_{ua_i}|$ the mean Pearson correlation between the utility and the attributes, and $|\bar{\rho}_{aia_j}|$ the mean intercorrelation between the attributes. CPU stands for central processing unit and f. lamp for fluorescent lamp.

Environment	No. A	No. a	R^2	$ max(\rho_{ua_i}) $	$ \bar{\rho}_{ua_i} $	$ \bar{\rho}_{aia_j} $
Beer aroma*	23	7	0.39	0.35	0.23	0.43
Cheese taste	30	3	0.65	0.76	0.67	0.65
CPU efficiency*	209	6	0.87	0.86	0.64	0.51
F. lamp lifetime*	14	3	0.57	0.33	0.24	0.41
Machine productivity	40	5	0.56	0.71	0.21	0.07
Octane quality	82	4	0.91	0.87	0.63	0.51
Olive oil quality	16	4	0.65	0.51	0.36	0.15
Potato taste	32	4	0.35	0.35	0.28	0.11
Red wine quality*	1599	11	0.36	0.48	0.19	0.2
Seed yield	30	3	0.70	0.81	0.36	0.18
Tea quality	64	5	0.39	0.43	0.21	0.21
White wine quality*	4898	11	0.28	0.44	0.15	0.17

dom from the initial dataset. This transformation was implemented to generate a plausible decision-making ecology. The ecological characteristics of the resulting environments were very similar to those of the complete environments reported in Table 1.⁴ We used the same subset of 200 alternatives consistently in all the analyses reported in this paper.

3.2 Performance

The performance of a model depends (1) on the order it generates, which determines how quickly it discovers high quality alternatives, and (2) on the estimated standard deviation of the error component $\hat{\sigma}$, which influences the subjective probability that the utility of the next alternative to be sampled will be higher than the utility of the best alternative discovered up to that point. Clearly, when $\hat{\sigma}$ increases, the subjective probability and consequently the expected returns from further search also increase. Finally, the model performance depends (3) on the cost of search in the environment. Factor 3 is an environmental factor, while factors 1 and 2 reflect how well a model captures the properties of the environment.

At first, we manipulated the cost of search and pitted the models against each other in eight different cost con-

ditions. For the purposes of our simulation we adapted the technique of cross-validation to a search problem. We fixed the parameters β_j and σ corresponding to each model in half of the data set (training set) and evaluated the performance of the models in the remaining half (test set). This process was repeated 10,000 times in total. For each repetition the alternatives that were part of the training and test sets were drawn at random from the entire data set. As a result, the maximum utility, which could be achieved by all models in each of the repetitions, differed slightly depending on the utilities of the alternatives that were part of the training and test sets. The models first sampled the alternative that they estimated to have the highest expected utility. Then each model decided according to equation 1 whether to proceed to search the alternative with the second highest expected utility, then the third, and so forth. When a model stopped search, the alternative with the highest utility among the alternatives searched up to that point was chosen. The measure of performance was the average utility achieved by a model in 10,000 repetitions. In real life, this corresponds to the average performance of 10,000 decision makers, with randomly sampled experiences, who all followed the optimal policy. This measure can also be seen as an approximation of the expected payoff for a single decision maker.

Average results from the 12 environments for all eight cost conditions are presented in Table 2 and the results from each of the 12 environments individually are pre-

⁴For red wine quality: $R^2 = 0.3$, $|max(\rho_{ua_i})| = 0.43$, $|\bar{\rho}_{ua_i}| = 0.16$, $|\bar{\rho}_{aia_j}| = 0.20$; for white wine quality: $R^2 = 0.35$, $|max(\rho_{ua_i})| = 0.48$, $|\bar{\rho}_{ua_i}| = 0.16$, $|\bar{\rho}_{aia_j}| = 0.19$.

Table 2: The average performance of the models across the 12 environments in 8 cost conditions. Multi-attribute linear utility (MLU) performed best but as the cost decreased performance differences with equal-weighted linear utility (EW) and single-attribute utility (SA) attenuated.

Cost	Average performance		
	MLU	EW	SA
1/2 ³	0.606	0.594	0.561
1/2 ⁴	0.715	0.696	0.681
1/2 ⁵	0.783	0.775	0.761
1/2 ⁶	0.824	0.822	0.816
1/2 ⁷	0.857	0.851	0.856
1/2 ⁸	0.885	0.878	0.881
1/2 ⁹	0.902	0.899	0.901
1/2 ¹⁰	0.912	0.912	0.911

sented in Figure 2. The best performing model was MLU. The performance differences were largest for high search costs and they gradually attenuated as the costs decreased. Clearly, a high cost implies that the models are likely to search only a few alternatives at the beginning of the order, and the models may have different orders. As the costs decrease, all the models search further down the expected utility line and are more likely to discover the high-quality alternatives in the environment. As a result, the performance differences diminish.

There is much variability in performance in individual environments, where in several cases the simpler models outperformed MLU for the entire cost range. For example, EW performed best for most of the cost conditions in the beer aroma, cheese taste, CPU efficiency, olive oil quality, and potato taste environments, while SA performed best in the red wine quality, tea quality, and octane quality environments. Thus, although MLU performed better on average, it clearly outperformed the simpler models only in the remaining four environments. The better performance of MLU at the aggregate level was mainly driven by its superior performance in the white wine quality environment. As stressed earlier, the model-specific factors that may influence the models' performance are on the one hand the order generated by each model and on the other hand the estimated standard deviation. In the rest of the Results section we decouple the roles of these two factors.

3.3 Impact of search order on performance and search length

To disentangle the impact of the search order from that of the estimated standard deviation of the error compo-

nent $\hat{\sigma}$, we first examined the average return to search (as measured by the average utility of the best searched alternative), achieved by each multi-attribute utility model, for all possible search lengths k , without implementing any search cost. Moreover, we compared these results to random search, which corresponds to the assumption made in most optimal-stopping studies. As we did for the full task, we cross-validated the models on half of the data points and we repeated our simulations 10,000 times.

As shown in Figure 3, the largest difference between the subjective models and random search is found in environments with high R^2 of the best fitting linear regression. In the environments with the highest R^2 , such as CPU efficiency and octane quality, the best solution was almost always located in one of the first search trials by all the multi-attribute models. In contrast, in environments with low R^2 such as beer aroma, potato taste, and fluorescent lamp lifetime, the margin is smaller and in some cases random search performed almost as well as the multi-attribute models.

In all environments there is a close correspondence between the average return to search as a function of the search length and the performance of the model in the full task, as depicted in Figure 2. For most costs, the model with the highest average return to search is also the best performing model in the full task. In four of the five environments in which EW performed best in the full task (beer aroma, cheese taste, CPU efficiency, potato taste) it also performed best for most of the search lengths. Similarly, in the three environments where SA performed best in the full task (red wine quality, tea quality and octane quality) it also performed as well as or better than the other two models for most search lengths in the search task. Overall, there are a few cases, such as the tea quality environment, where different models lead to best performance for different search lengths k .

Note that the factors that influence the performance of a model directly influence the length of search, as observed in Figure 4. According to the theory presented in section 2, the utility of the best alternative secured so far, y , and the probability $P(u_k > y)$, that the utility u_k of the next alternative down the expected utility line A_k will be higher than y should be inversely related. This suggests that models that place high-utility alternatives earlier in the search order will be characterized by lower returns to search and should, ceteris paribus, search less. Indeed, this is what we observe in Figures 3 and 4. However, to fully understand the performance and search length of the models we also need to come to grips with the role played by the standard deviation of the error component σ . The next section delves into that exactly.

Figure 2: The average utility achieved by the subjective utility models as a function of search cost. Overall, multi-attribute linear utility (MLU) performed best, but equal-weighted linear utility (EW) and single-attribute utility (SA) also had regions of superior performance. The difference between the models gets smaller as the cost of search decreases. CPU stands for central processing unit and f. lamp for fluorescent lamp.

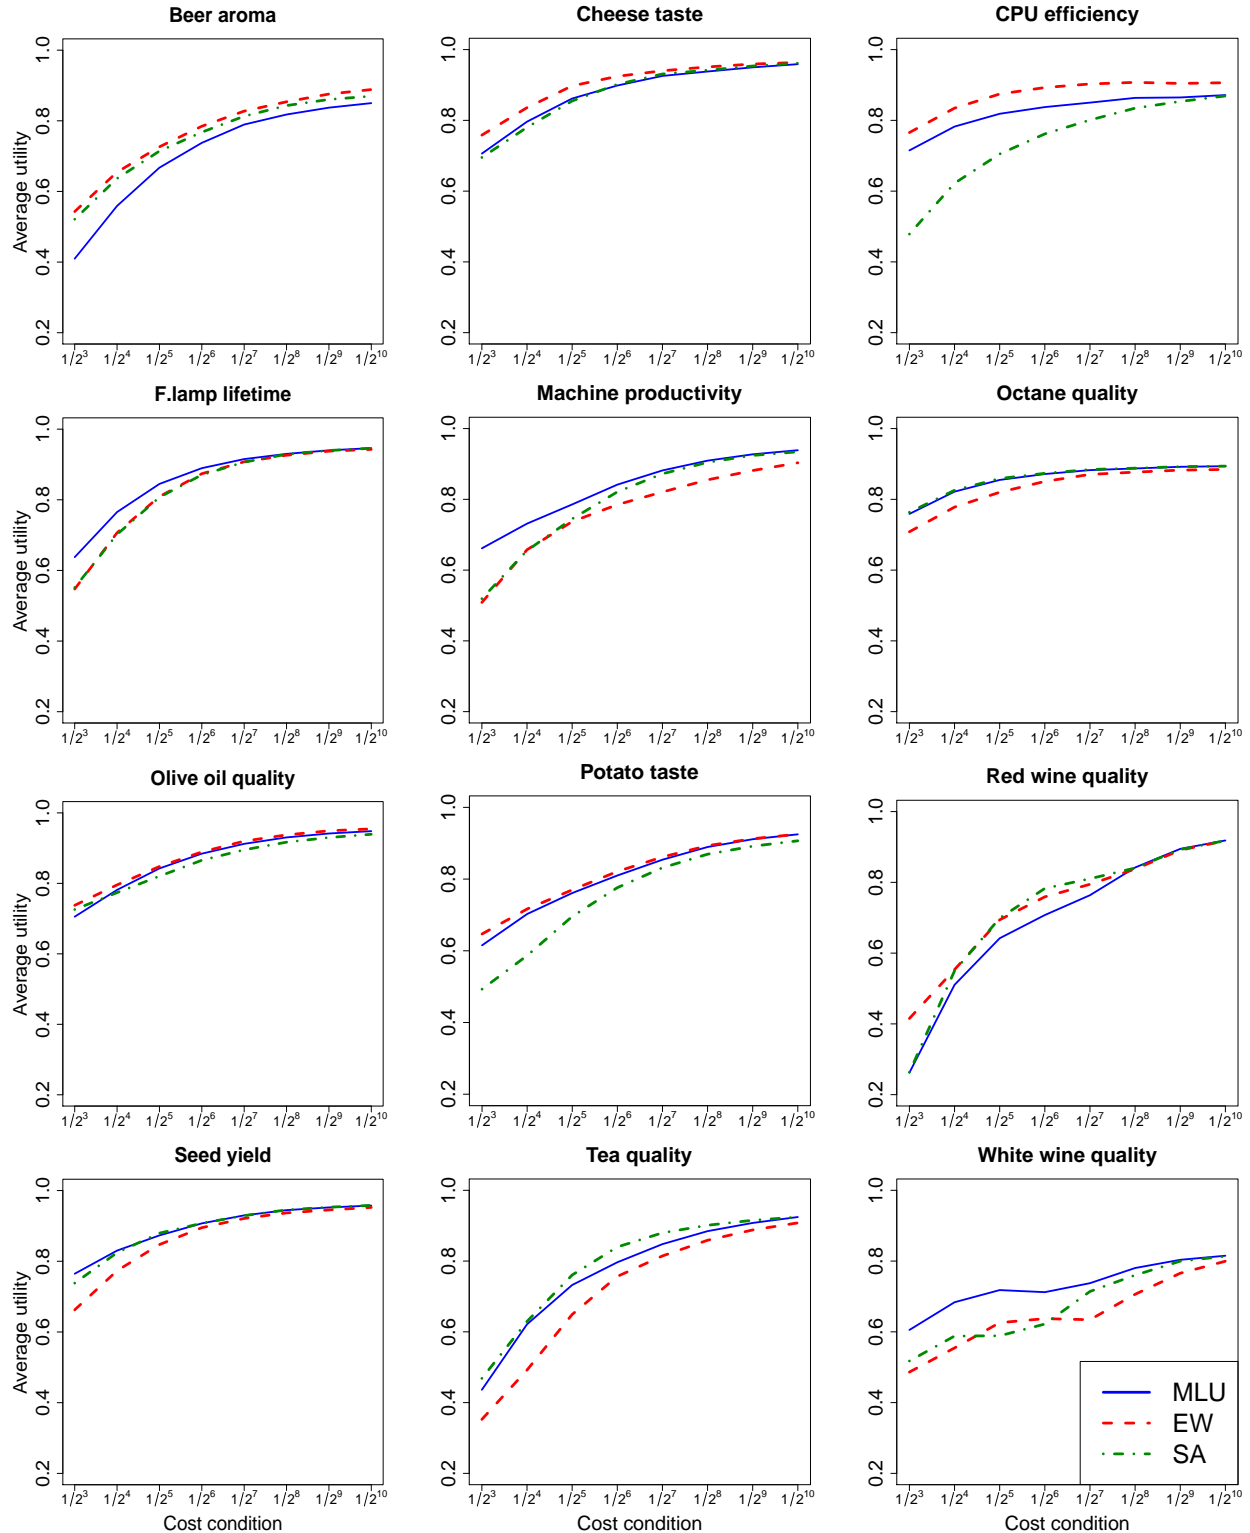


Figure 3: The average utility of the best explored alternative after a search of length k . In environments with a high R^2 , such as central processing unit (CPU) efficiency and octane quality, the best alternatives are located early in the search and the differences between these strategies and random search are the largest. In most cases there is a close correspondence between model performance in the full task and model performance in mere search; contrast with Figure 2. MLU stands for multi-attribute linear utility, EW for equal-weighted linear utility, and SA for single-attribute utility.




Figure 4: On the upper left side we present the average standard deviation of the error component of each of the models in that environment. The length of search is moderated by the best alternative discovered so far and the estimated deviation of the error component of the models. On average, multi-attribute linear utility (MLU) searches less than the equal-weighted linear utility (EW) and single-attribute utility (SA) models but there is some variability across environments.

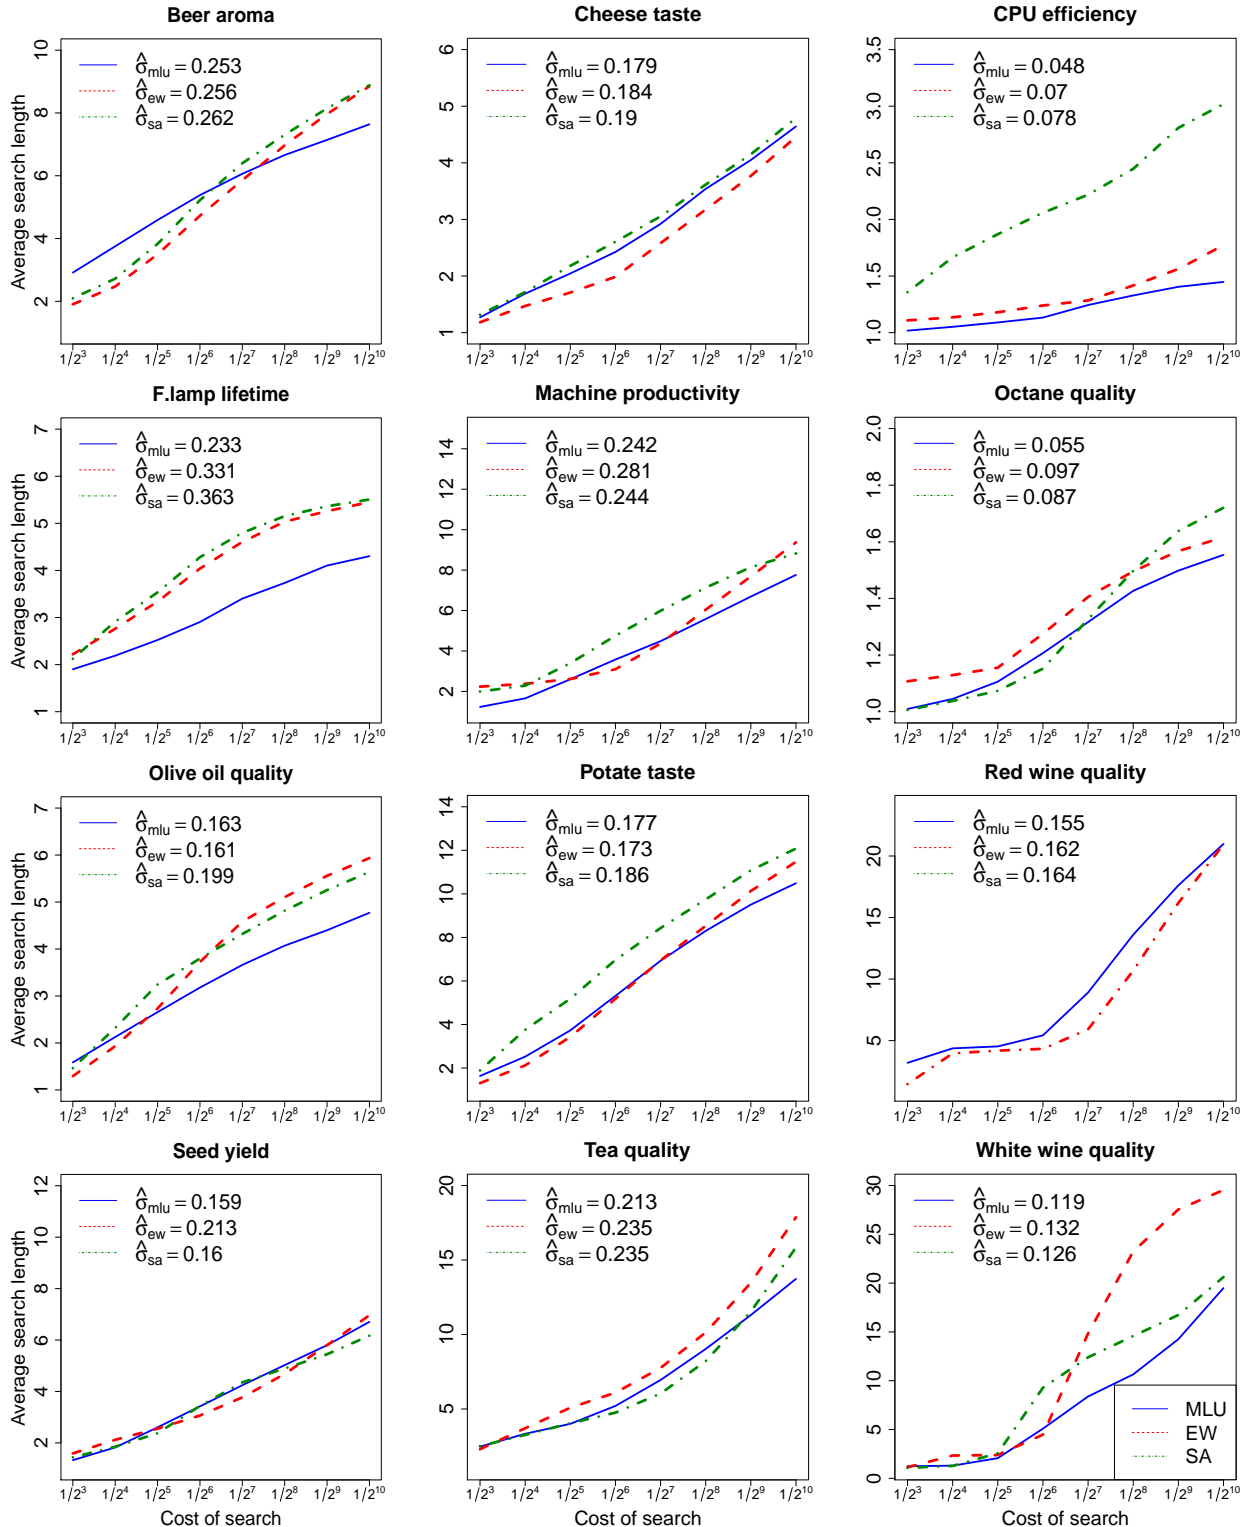


Table 3: The average length of search (number of alternatives sampled k) across environments for the three models for different costs of search. The equal-weighted linear utility (EW) and single-attribute utility (SA) models search more alternatives, on average, than multi-attribute linear utility (MLU). The differences in the length of search are more pronounced for low costs.

Cost	Average search length		
	MLU	EW	SA
$1/2^3$	1.74	1.58	1.74
$1/2^4$	2.23	2.28	2.45
$1/2^5$	2.81	2.83	3.21
$1/2^6$	3.70	3.62	4.47
$1/2^7$	4.88	5.39	5.58
$1/2^8$	6.11	7.24	6.72
$1/2^9$	7.34	8.85	8.08
$1/2^{10}$	8.66	10.45	9.54

3.4 The role of the standard deviation σ

The second factor that has an impact on the performance of the model and the length of search is the standard deviation of the error component σ . To assure good performance the models' error component should correspond to the unexplained uncertainty in the environment. For estimated $\hat{\sigma} = 0$ the model would deem that the first alternative is also the best one. For $\hat{\sigma} = \infty$ the model would calculate $P(u_k > y) = 0.5$ (following equation 1). However, both these models would be unadaptive in environments in which the uncertainty that cannot be explained by $f(\mathbf{a}_i)$ is moderate. A lower standard deviation of the error component $\hat{\sigma}$ implies a decrease in the subjective probability $P(u_k > y)$ and a decrease in the expected returns from finding a better alternative $E(u_k - y | u_k > y)$. Thus, ceteris paribus, it should lead to a shorter search.

Within models, the estimated standard deviation of the error component $\hat{\sigma}$ in our simulations is inversely related to the accuracy of the estimates of the model in the training set in which the model was fitted. Consequently, in most cases a good ordering should be accompanied by a relatively low error component and in tandem they lead to a reduced length of search. Note, however, that there could be a discrepancy between the accuracy of the estimates of a model and the average returns to search. Remember that only the best alternative discovered so far counts for the decision maker in the search task.

Between models, the MLU tends to capture a larger proportion of the uncertainty in \mathbf{a}_i ($\hat{\sigma}_{mlu} = 0.166$) and usually has a slightly lower estimated standard deviation $\hat{\sigma}$

than EW and SA ($\hat{\sigma}_{ew} = 0.191$ and $\hat{\sigma}_{sa} = 0.191$). As reported in Figure 4, only in the olive oil quality and potato taste environments was a heuristic model (EW) found to have an estimated standard deviation of the error component lower than the full model (MLU).

As seen in Figures 2-4, the models that achieve higher average returns to search (Figure 3), search less (Figure 4, Table 3 for average results) and score better on the full task (Figure 2). The differences in the estimated standard deviation of the models $\hat{\sigma}$ are in most environments marginal. In half of the environments they correspond to performance differences in the average returns to search. This suggests that the error components of the models are sufficiently well calibrated to the unexplained uncertainty in the environment. There are a few cases where a different estimated standard deviation $\hat{\sigma}$ might lead the models to markedly different performance. For example, in the CPU efficiency environment the EW and MLU have similar returns to search. MLU has a much lower estimated standard deviation $\hat{\sigma}$, searches less than EW, and performs worse in the full task. This indicates that MLU may occasionally miss a very good alternative and it could benefit from searching more.

3.5 Paired-comparison results

So far, we have discussed how different properties of the search models influence their performance and we have looked at specific environments to understand how these factors play out in practice. We have illustrated that the most crucial factor is the search order in which the models sample the alternatives, although this has to be accompanied by a well-calibrated estimated standard deviation of the error component $\hat{\sigma}$. But is there a way to predict which strategy is most likely to be successful in a given environment? The conditions under which different linear or heuristic models perform well in choice and inference contexts have been thoroughly investigated (for a review see Katsikopoulos, 2011). Few studies generalize beyond binary choice to choice between several alternatives or among the entire data set; however, most of the existing literature has focused on binary choices. To answer our question we attempt to shed light on the connections between the novel task we presented and the well-studied paired-comparison task.

Thus, we examined whether there is a correspondence in performance between the binary choice and the search task. We compared the performance of the three models in the binary choice task. As before, we fixed the parameters corresponding to each model in half of the dataset and evaluated their performance in all possible binary choice tasks in the remaining half. The process was repeated in total 10,000 times.

As seen in Table 4, on average, MLU performed best, followed by SA and then EW. In the individual environ-

Table 4: Accuracy of the three models on binary choices. Fifty percent of the dataset was used as a training set and all the possible pairs of the test set were used to evaluate the models. Multi-attribute linear utility (MLU) performed best, on average, followed by single-attribute utility (SA) and equal-weighted linear utility (EW). MLU performed best in seven environments EW in four and SA in one.

Environment	Subjective models		
	MLU	EW	SA
Beer aroma	0.48	0.49	0.47
Cheese taste	0.79	0.79	0.76
CPU efficiency	0.83	0.85	0.80
F. lamp lifetime	0.67	0.54	0.55
Machine productivity	0.64	0.55	0.63
Octane quality	0.86	0.72	0.79
Olive oil quality	0.72	0.71	0.58
Potato taste	0.63	0.64	0.57
Red wine quality	0.63	0.60	0.63
Seed yield	0.78	0.64	0.77
Tea quality	0.69	0.62	0.63
White wine quality	0.65	0.60	0.66
Mean performance	0.70	0.65	0.65

ments, the performance of the models varied significantly. There were four environments in which EW performed best and one in which SA did. In general, the performance in the binary choice task is a good proxy of performance in the full search task. The model that performed best in the paired-comparison task also performed best in 7 of the 12 environments in the search task for cost equal to $1/2^3$. The discrepancies observed can be attributed to the fact that in the full search task, the alternatives that are searched early on contribute disproportionately to the success of a model. In contrast, in binary choices all possible single choices in the data set contribute equally to the performance of the model. Hence, it is possible that the model that performs best in the search task does not maintain that superior performance in the binary choice task, and vice versa.

4 Discussion

4.1 Conceptual implications of our framework

4.1.1 Deterministic optimization and search models: A possible compromise

In optimization problems, widely studied in economics, decision makers can determine the alternative that maximizes their utility. This vision of decision making con-

trasts with search models in which the decision makers sample alternatives at random and stop search after encountering a good enough alternative (e.g., Simon, 1955; Chow, Robbins & Siegmund, 1971; Caplin, Dean & Martin, 2011). Random sampling may lead to violations of the revealed preference principle and unpredictability in regard to the choices of individual decision makers. Ordered search models provide a possible compromise between these two approaches. Decision makers have a well-defined utility function before the search starts. However, as long as there is some uncertainty about the exact utility of the alternatives, it may pay to sample some of them to learn their utility. In our model, the initial preferences guide the search process but are also subject to revision when the true utility of the sampled alternatives is revealed. For an external observer, such as a firm or a market analyst, ordered search is more predictable, at the level of an individual decision maker, than random search. In ordered search, if the model of the decision maker and the actual utility of the alternatives are known, the external observer could also predict the final decisions made, as well as the preference reversals that would occur along the way.

4.1.2 Evaluating the hypotheses of ordered search models

Employing our framework, we can examine the conditions under which the assumptions postulated in previous ordered search models hold true. Bagwell and Ramey (1994) suggested that decision makers may use a simple rule of thumb and first consider buying products from firms that advertise more.⁵ This corresponds to an SA model where the amount of advertising is the most informative cue. This, indeed, is plausible in some cases. However, the SA model implies that when more informative attributes are available, the amount of advertisement may be completely ignored as an attribute. Armstrong et al. (2009) assumed that consumers search products according to an abstract attribute called prominence and suggested that firms might be willing to invest to achieve prominence and improve their search order. Our framework allows us to estimate the exact impact of an intervention in the attributes of a product in terms of the firm’s prominence in the market. Finally, Moorthy et al. (1997) suggested that the expectation of the utility of a brand’s products is normally distributed. In addition, they suggested that more experienced consumers are better able to differentiate between products. These assumptions are both in line with our framework.

⁵A similar argument for the role of advertisement in the domain of choice has been made by Goldstein and Gigerenzer (2002).

4.2 Applied implications of our framework

4.2.1 Sequential search and consideration set formation

Our modeling approach suggests that only the alternatives that have been sampled by the decision makers stand a chance of being selected. Similarly, several marketing scientists have advanced choice models in which the decision makers first restrict their attention to a subset of the alternative set — commonly called the consideration set. The decision makers then examine the alternatives of this subset more closely and finally choose one of the alternatives in it (e.g., Wright & Barbour, 1977; Shocker, Ben-Akiva, Boccara & Nedungadi, 1991; Gilbride & Allenby, 2004; Moe, 2006). Such models have often been found to outperform, in fitting and prediction, discrete choice models in which the decision makers are assumed to consider all the alternatives (e.g., Gilbride & Allenby, 2004). Our approach shares some of its assumptions with a popular model of consideration set formation put forward by Roberts and Lattin (1991). In their model, decision makers, whose pay-off function is described by a compensatory multi-attribute utility model with an additional error term, decide which alternatives to include in their consideration set. Similar to our approach, the decision makers examine the alternatives in the order of decreasing expected utilities predicted by their utility model, paying a fixed cost for every new alternative they place in their consideration set. In contrast to our model, the decision makers do not learn the exact utility of the alternatives immediately after paying the cost. Instead, they learn all the utilities of the alternatives in the consideration set right before they choose among them. Our model can be seen as the sequential search counterpart to Roberts and Lattin theory of consideration set formation. Inversely, Roberts and Lattin's model can also be understood as an ordered search model in which the search length has to be decided at the outset. The exact domain of application of sequential search and fixed-sample-size models of consideration-set formation may depend on the exact characteristics of the decision making context and it should be the subject of further empirical investigation in the future.

4.2.2 Parallels to online ranking schemes

There are clear-cut parallels between the search approach we present here and the methods used by commercial search engines and recommendation systems to pre-rank the alternatives for Internet users. Such ranking schemes implement a valence function that maps the attributes of the alternatives to the relevance or utility for the user (Burges, Shaked, Renshaw, Lazier, Ari and Deeds, Matt, Hamilton & Hullender, 2005; Hüllermeier, Fürnkranz & Cheng, Weiwei and Brinker, Klaus, 2008; Yaman, Walsh,

Littman & Desjardins, 2011). The valence functions underlying the ranking schemes are trained with clickstream decision data or other information that can reveal the preferences of the user (Sarwar, Karypis, Konstan & Riedl, 2000). Then, similar to what our theory suggests, the ranking schemes present the alternatives in decreasing order of relevance or utility. This approach is implemented without invoking a formal decision-making theory that predicts how decision makers will choose on the basis of the presented rank order. If the goal of the ranking-scheme engineers is to increase the utility derived by the users, a formal decision-making theory, such as search theory, might further inform the development of ranking-schemes as well as the techniques used to train their valence functions. Indeed, there have been recent papers in machine learning about why the design of ranking techniques can benefit from taking into account how people actually decide in rank-ordered environments (Agichtein, Brill & Dumais, 2006; Chapelle, Metzler, Zhang & Grinspan, 2009). We take a step in that direction and illustrate the role of cost of search and the uncertainty in the environment in the search process.

4.3 Connections to prescriptive and descriptive decision making

4.3.1 From choice to search

For the paired-comparison problem, where the task is to choose one of two alternatives, the accuracy of heuristics such as EW and SA has been analyzed and compared to that of the full linear model in numerous studies. Until now, a thread of the existing literature in judgment and decision making has examined binary choices in environments with binary or continuous attributes without making any assumptions about the mapping from an alternative's attributes to its utility (e.g., Katsikopoulos et al., 2010) or when the mapping is characterized by noise (Hogarth & Karelaia, 2005; Hogarth & Karelaia, 2006; Rieskamp & Otto, 2006; Davis-Stober, Dana & Budescu, 2010a; Davis-Stober et al., 2010b). A second thread of the literature has focused on environments with binary or continuous attributes, where a mapping between the attributes and the utility exists but decision makers have imprecise knowledge about the attribute weights (Johnson & Payne, 1985; Martignon & Hoffrage, 2002; Hogarth & Karelaia, 2005; Baucells, Carrasco & Hogarth, 2008; Katsikopoulos, 2013).

Given the observed correspondence in choice task and search task results, the findings of the first thread of studies in binary choice tasks may also generalize to the search task. Overall, these studies have found no large performance differences between the heuristics and the full model and that heuristics can outperform the full model

under the appropriate conditions. Both EW and SA fare especially well in out-of-sample prediction (Einhorn & Hogarth, 1975; Hogarth & Karelaia, 2005; Katsikopoulos et al., 2010). The SA model tends to perform well when a simply or cumulatively dominating alternative is present (Baucells, Carrasco & Hogarth, 2008; Şimşek, 2013; Katsikopoulos, Egozcue & Garcia, 2014), or when there exist high correlations between the single attribute and all other attributes (Hogarth & Karelaia, 2005; Davis-Stober et al., 2010a, 2010b). EW tends to perform well when the variability in cue validities is small or when there are high intercorrelations between all the attributes (Einhorn & Hogarth, 1975; Wainer, 1976). In the environments that we studied we also found support for some of these findings. For example in the four environments, in which EW performed best in binary choice the difference between $|max(\rho_{ua_i})|$ and $|\bar{\rho}_{ua_i}|$ was small.

Although our first results suggest a relation between the search task and the choice task, additional research is required in the future to establish this relation and to identify the conditions under which differences in the relative performance of models in the two tasks are to be expected. A further follow-up to our study would be to examine when decision makers choose a certain strategy (Bröder, 2003) and how the decision makers learn to act adaptively and to select a strategy that performs well in a given environment (e.g., Rieskamp & Otto, 2006).

4.3.2 Psychological plausibility of the stopping rule

So far we have shown that the proposed stopping policy is optimal. But is it psychologically plausible? When sampling from a known distribution with or without recall an optimally acting decision maker should always stop right after encountering an alternative with a value higher than the optimal threshold. Several variations of such optimal threshold problems have been studied extensively in psychology and economics. We have identified five studies reporting results on experiments in which subjects sampled from a known distribution with recall (Rapoport & Tversky, 1970; Schotter & Braunstein, 1981; Hey, 1987; Kogut, 1990; Sonnemans, 1998). In sum, a moderate proportion of the participants in these studies behaved in a manner consistent with the optimal stopping rule. Common discrepancies from the optimal strategy included stopping too early and exercising recall. Nonetheless, the researchers found that the average performance of the participants was near optimal. Hey (1982) and Sonnemans (1998) reported that many subjects used heuristic strategies that appeared consistent with the optimal rule and led to near optimal performance. In the stopping policy we have presented, the decision maker should at every search step reevaluate the returns from sampling the next alternative in line. This task may appear demanding in

relation to the optimal threshold rule. However, it has the same structure as a simplified version of a signal-detection problem (e.g., Green & Swets, 1966) — in which humans are known to perform fairly well. Thus, we believe that the stopping policy could be psychologically plausible as such, or it could be well approximated by clever heuristic algorithms. Clearly, human behavior in the task has to be investigated experimentally in the near future.

4.3.3 Alternative search algorithms

We showed that, under the homoscedasticity assumption embedded in many linear models and the subjective linear utility models outlined in this paper, the intuitive policy of searching the alternatives in the order of their subjective expected utility is optimal. In addition, in this special case our policy coincides with another simple algorithm called *directed cognition*, which has been proposed by Gabaix, Laibson, Moloche and Weinberg (2006). This algorithm searches myopically, as if the next step of the search process was also the last one. When the assumption of the homogeneously distributed error component ϵ does not hold, the behavior suggested by our ordering policy diverges from Weitzman's (1979) optimal solution and from Gabaix et al.'s (2006) directed cognition algorithm. Gabaix et al. (2006) showed that, even in simple cases of indexable problems like those discussed by Weitzman (1979) and Gittins, Glazebrook and Weber (1989), the decision makers are unlikely to follow the optimal algorithm and instead decide in line with the directed cognition algorithm. In the future, as in Gabaix et al. (2006), one could examine experimental scenarios in which the behavior prescribed by the three discussed algorithms diverges to evaluate their psychological plausibility in different decision-making environments.

4.4 Extensions and limitations

The current model could be extended to scenarios with different search costs for different alternatives or when the decision maker can choose an alternative directly without paying the sampling cost. Further, the model could be readily extended to cases when the decision maker can choose more than one alternative. In addition, it is possible to account for contexts where new alternatives become available at a later point in time. The decision maker can construct a search order for the new alternatives and consider if it is worthwhile to sample some of them, examining first the alternative with the highest expected utility.

In our model we have assumed that decision makers have direct and free access to the attributes all at once. We did not discuss cases where the decision makers have to pay a search cost to learn additional attributes before moving forward to examine further alternatives. A model

of this kind with random sampling is presented by Lim, Bearden and Smith (2006) and has been investigated empirically by Bearden and Connolly (2007). Further, as in search problems, we have assumed that the decision maker learns the exact utility of an alternative after paying a search cost and examining it. However, there are many dynamic decision-making contexts where the cost is internally defined as an opportunity cost when consuming an inferior alternative (Nelson, 1970). In these environments information acquisition can be inherently noisy and decision makers may want to sample the alternatives repeatedly. Such decision-making contexts are commonly referred to as multi-armed bandits. In fact, when we change the costly sampling assumption to repeated sampling, in which the final pay-off is the sum of the experienced utilities, our framework turns into a multi-armed bandit with contextual information. This type of bandit framework has been receiving increasing attention in recent years in machine learning (Pavlidis, Tasoulis & Hand, 2008; Li, Chu, Langford & Schapire, 2010).

So far we have postulated that decision makers have stable preferences and an accurate error estimate of their model throughout the entire search process. This strong assumption may hold precisely or it may approximate the truth in many decision-making environments, especially when decision makers have long experience testing alternatives. However, in some environments decision makers may not know the utility weights but rather learn them along the way as they examine new alternatives. This approach has been followed by Dzyabura (2013), who assumed that decision makers update the estimates of the weights and the search order after each new alternative they examine. Clearly, when decision makers learn their preferences the ordering and stopping rules derived in our paper are not guaranteed to be optimal. A fully rational policy would have to anticipate the future evolution of decision makers' preferences and then build any beliefs about preference change into the search and stopping rules. Even in simplified scenarios this approach is known to be computationally intractable. In similar preference-learning problems encountered in machine learning, greedy heuristics are implemented instead to balance preference learning and exploitation (Brochu, De Freitas & Ghosh, 2007; Brochu, Cora & De Freitas, 2010). Another approach would be to compare, for any given search length, the performance of alternative, optimal in some respect, active learning algorithms (Fedorov, 1972; Sugiyama & Nakajima, 2009).

4.5 Conclusion

In two recent publications Luan, Schooler and Gigerenzer (2011, 2014) stressed the need to integrate decision-making theories in psychology and illustrate how appar-

ently disparate models share common conceptual ground. In the same vein, we argued that choice and search problems, which until now have been studied separately, are the boundary cases of a broader decision-making problem. We showed how three choice models that have been extensively studied in the field of judgment and decision making can guide the search for good alternatives and we formulated their corresponding optimal stopping rule. Then, we compared the performance of the models in 12 real-world environments ranging from consumer choice to industrial experimentation and illustrated how each models' expected utility ordering and estimation error influence its' performance and length of search. As in previous model comparisons, in one-shot choice problems we found that heuristic linear models performed on average close to a multi-attribute linear utility model. Moreover, in individual environments the heuristic models often outperformed the full model. To further understand when such results are to be expected we examined the relationship between the search problem and the well-studied binary choice problem. We found that in most cases the models that performed well on the binary choice task also did so in the search task. This suggests that previous findings on the ecological rationality of choice and inference strategies are also relevant to the search task. Finally, we discussed the connections of our model to the existing literature and suggested possible paths for future research.

References

- Agichtein, E., Brill, E. & Dumais, S. (2006). Improving web search ranking by incorporating user behavior information. *Proceedings of the 29th annual international ACM SIGIR conference on research and development in information retrieval* (pp. 19–26).
- Arbatskaya, M. (2007). Ordered search. *The RAND Journal of Economics*, 38(1), 119–126.
- Armstrong, M., Vickers, J. & Zhou, J. (2009). Prominence and consumer search. *The RAND Journal of Economics*, 40(2), 209–233.
- Bagwell, K. & Ramey, G. (1994). Coordination economies, advertising, and search behavior in retail markets. *The American Economic Review*, 84(3), 498–517.
- Barron, F. H. (1987). Influence of missing attributes on selecting a best multiattribute alternative. *Decision Sciences*, 18(2), 194–205.
- Baucells, M., Carrasco, J. A. & Hogarth, R. M. (2008). Cumulative dominance and heuristic performance in binary multiattribute choice. *Operations Research*, 56(5), 1289–1304.
- Bearden, J. N. & Connolly, T. (2007). Multi-attribute sequential search. *Organizational Behavior and Human*

- Decision Processes*, 103(1), 147–158.
- Brochu, E., Cora, V. M. & De Freitas, N. (2010). A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*.
- Brochu, E., De Freitas, N. & Ghosh, A. (2007). Active preference learning with discrete choice data. *Advances in neural information processing systems* (pp. 409–416).
- Bröder, A. (2003). Decision making with the "adaptive toolbox": Influence of environmental structure, intelligence, and working memory load. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 29(4), 611–625.
- Burges, C., Shaked, T., Renshaw, E., et al., Hamilton, N. & Hullender, G. (2005). Learning to rank using gradient descent. *Proceedings of the 22nd international conference on Machine learning* (pp. 89–96).
- Caplin, A., Dean, M. & Martin, D. (2011). Search and satisficing. *The American Economic Review*, 101(7), 2899–2922.
- Chapelle, O., Metlzer, D., Zhang, Y. & Grinspan, P. (2009). Expected reciprocal rank for graded relevance. *Proceedings of the 18th ACM conference on Information and knowledge management* (pp. 621–630).
- Chow, Y. S., Robbins, H. & Siegmund, D. (1971). *Great expectations: The theory of optimal stopping*. Houghton Mifflin.
- Davis-Stober, C. P., Dana, J. & Budescu, D. V. (2010). A constrained linear estimator for multiple regression. *Psychometrika*, 75(3), 521–541.
- Davis-Stober, C. P., Dana, J. & Budescu, D. V. (2010). Why recognition is rational: Optimality results on single-variable decision rules. *Judgment and Decision Making*, 5(4), 216–229.
- Dawes, R. M. & Corrigan, B. (1974). Linear models in decision making. *Psychological Bulletin*, 81(2), 95.
- DeGroot, M. (1970). *Optimal statistical decisions*. New York [etc.]: McGraw-Hill.
- Dzyabura, D. (2013). The role of changing utility in product search. Available at SSRN: <http://ssrn.com/abstract=2202904>.
- Einhorn, H. J. & Hogarth, R. M. (1975). Unit weighting schemes for decision making. *Organizational Behavior and Human Performance*, 13(2), 171–192.
- Fasolo, B., McClelland, G. H. & Todd, P. M. (2007). Escaping the tyranny of choice: When fewer attributes make choice easier. *Marketing Theory*, 7(1), 13–26.
- Fedorov, V. V. (1972). *Theory of optimal experiments*. Academic Press Inc.
- Gabaix, X., Laibson, D., Moloche, G. & Weinberg, S. (2006). Costly information acquisition: Experimental analysis of a boundedly rational model. *The American Economic Review*, 96(4), 1043–1068.
- Gigerenzer, G. & Goldstein, D. G. (1996). Reasoning the fast and frugal way: Models of bounded rationality. *Psychological review*, 103(4), 650–669.
- Gigerenzer, G., Todd, P. M. & et al. (1999). *Simple heuristics that make us smart*. Oxford University Press, USA.
- Gilbride, T. J. & Allenby, G. M. (2004). A choice model with conjunctive, disjunctive, and compensatory screening rules. *Marketing Science*, 23(3), 391–406.
- Gittins, J., Glazebrook, K. & Weber, R. (1989). *Front matter*. Wiley Online Library.
- Goldstein, D. G. & Gigerenzer, G. (2002). Models of ecological rationality: The recognition heuristic. *Psychological Review*, 109(1), 75–90.
- Green, D. M. & Swets, J. A. (1966). *Signal detection theory and psychophysics*. Wiley New York.
- Hammond, K. R., Hursch, C. J. & Todd, F. J. (1964). Analyzing the components of clinical inference. *Psychological Review*, 71(6), 438–456.
- Hey, J. D. (1987). Still searching. *Journal of Economic Behavior & Organization*, 8(1), 137–144.
- Hey, J. D. (1982). Search for rules for search. *Journal of Economic Behavior & Organization*, 3(1), 65–81.
- Hogarth, R. M. & Karelaia, N. (2005). Ignoring information in binary choice with continuous variables: When is less “more”? *Journal of Mathematical Psychology*, 49(2), 115–124.
- Hogarth, R. M. & Karelaia, N. (2005). Simple models for multiattribute choice with many alternatives: When it does and does not pay to face trade-offs with binary attributes. *Management Science*, 51(12), 1860–1872.
- Hogarth, R. M. & Karelaia, N. (2006). Regions of rationality: Maps for bounded agents. *Decision Analysis*, 3(3), 124–144.
- Hüllermeier, E., Fürnkranz, J. & et al. (2008). Label ranking by learning pairwise preferences. *Artificial Intelligence*, 172(16), 1897–1916.
- Johnson, E. J. & Payne, J. W. (1985). Effort and accuracy in choice. *Management Science*, 31(4), 395–414.
- Katsikopoulos, K. V. (2013). Why Do Simple Heuristics Perform Well in Choices with Binary Attributes? *Decision Analysis*, 10(4), 327–340.
- Katsikopoulos, K. V., Schooler, L. J. & Hertwig, R. (2010). The robust beauty of ordinary information. *Psychological Review*, 117(4), 1259–1266.
- Katsikopoulos, K. V., Egozcue, M. & Garcia, L. F. (2014). Cumulative dominance in multi-attribute choice: Benefits and limits. *EURO Journal on Decision Processes*, 2, 1–11.
- Katsikopoulos, K. V. (2011). Psychological heuristics for making inferences: Definition, performance, and the emerging theory and practice. *Decision Analysis*, 8(1), 10–29.

- Keeney, R. L. & Raiffa, H. (1993). *Decisions with multiple objectives: Preferences and value trade-offs*. Cambridge University Press.
- Kogut, C. A. (1990). Consumer search behavior and sunk costs. *Journal of Economic Behavior & Organization*, 14(3), 381–392.
- Kohli, R. & Jedidi, K. (2007). Representation and inference of lexicographic preference models and their variants. *Marketing Science*, 26(3), 380–399.
- Lee, M. D. (2006). A Hierarchical Bayesian Model of Human Decision-Making on an Optimal Stopping Problem. *Cognitive Science*, 30(3), 1–26.
- Li, L., Chu, W., Langford, J. & Schapire, R. E. (2010). A contextual-bandit approach to personalized news article recommendation. *Proceedings of the 19th international conference on World Wide Web* (pp. 661–670).
- Lim, C., Bearden, J. N. & Smith, J. C. (2006). Sequential search with multiattribute options. *Decision Analysis*, 3(1), 3–15.
- Lippman, S. A. & McCall, J. (1976). The economics of job search: A survey. *Economic Inquiry*, 14(2), 155–189.
- Luan, S., Schooler, L. & Gigerenzer, G. (2014). From perception to preference and on to inference: An approach-avoidance analysis of thresholds. *Psychological Review*, 121(3), 501–525.
- Luan, S., Schooler, L. J. & Gigerenzer, G. (2011). A signal-detection analysis of fast-and-frugal trees. *Psychological Review*, 118(2), 316.
- Martignon, L. & Hoffrage, U. (2002). Fast, frugal, and fit: Simple heuristics for paired comparison. *Theory and Decision*, 52(1), 29–71.
- Moe, W. W. (2006). An empirical two-stage choice model with varying decision rules applied to internet click-stream data. *Journal of Marketing Research*, 43(4), 680–692.
- Moorthy, S., Ratchford, B. T. & Talukdar, D. (1997). Consumer information search revisited: Theory and empirical analysis. *Journal of Consumer Research*, 23(4), 263–277.
- Morgan, P. & Manning, R. (1985). Optimal search. *Econometrica: Journal of the Econometric Society*, 53(4), 923–944.
- Nelson, P. (1970). Information and consumer behavior. *The Journal of Political Economy*, 78(2), 311–329.
- Pavlidis, N. G., Tasoulis, D. K. & Hand, D. J. (2008). Simulation Studies of Multi-armed Bandits with Covariates. *Proceedings of the tenth international conference on computer modeling and simulation* (pp. 493–498).
- Payne, J. W., Bettman, J. R. & Johnson, E. J. (1993). *The adaptive decision maker*. Cambridge University Press.
- Rapoport, A. & Tversky, A. (1966). Cost and accessibility of offers as determinants of optional stopping. *Psychonomic Science*, 4(4), 145–146.
- Rapoport, A. & Tversky, A. (1970). Choice behavior in an optional stopping task. *Organizational Behavior and Human Performance*, 5(2), 105–120.
- Rieskamp, J. & Otto, P. E. (2006). SSL: A Theory of How People Learn to Select Strategies. *Journal of Experimental Psychology: General*, 135, 207–236.
- Roberts, J. H. & Lattin, J. M. (1991). Development and testing of a model of consideration set composition. *Journal of Marketing Research*, 28(4), 429–440.
- Sarwar, B., Karypis, G., Konstan, J. & Riedl, J. (2000). Analysis of recommendation algorithms for e-commerce. *Proceedings of the 2nd ACM conference on electronic commerce* (pp. 158–167).
- Schotter, A. & Braunstein, Y. M. (1981). Economic search: An experimental study. *Economic Inquiry*, 19(1), 1–25.
- Shocker, A. D., Ben-Akiva, M., Boccara, B. & Nedungadi, P. (1991). Consideration set influences on consumer decision-making and choice: Issues, models, and suggestions. *Marketing letters*, 2(3), 181–197.
- Simon, H. A. (1955). A behavioral model of rational choice. *The Quarterly Journal of Economics*, 69(1), 99–118.
- Şimşek, Ö. (2013). Linear decision rule as aspiration for simple decision heuristics. *Advances in Neural Information Processing Systems* (pp. 2904–2912).
- Sonnemans, J. (1998). Strategies of search. *Journal of Economic Behavior & Organization*, 35(3), 309–332.
- Stigler, G. J. (1961). The economics of information. *The Journal of Political Economy*, 69(3), 213–225.
- Sugiyama, M. & Nakajima, S. (2009). Pool-based active learning in approximate linear regression. *Machine Learning*, 75(3), 249–274.
- Wainer, H. (1976). Estimating coefficients in linear models: It don't make no nevermind. *Psychological Bulletin*, 83(2), 213.
- Weitzman, M. L. (1979). Optimal search for the best alternative. *Econometrica: Journal of the Econometric Society*, 47(3), 641–654.
- Wright, P. & Barbour, F. (1977). Phased decision strategies: Sequels to an initial screening. In *Multiple criteria decision making: North holland tims studies in the management sciences* (pp. 91–109). Amsterdam: North-Holland.
- Yaman, F., Walsh, T. J., Littman, M. L. & Desjardins, M. (2011). Democratic approximation of lexicographic preference models. *Artificial Intelligence*, 175(7), 1290–1307.
- Zwick, R., Rapoport, A., Lo, A. K. C. & Muthukrishnan, A. (2003). Consumer sequential search: Not enough or too much? *Marketing Science*, 22(4), 503–519.

5 Appendix 1: Proof of result 1

First notice that for any alternative A_k

$$R(A_k) = P(u_k > y) \times E(u_k - y | u_k > y) - c$$

$$= \int_y^\infty (u - y) dF_K(u) - c = \int_0^\infty u dF_K(u) - c$$

For alternative A_i and A_j let $\alpha = f(\mathbf{a}_i) - f(\mathbf{a}_j)$. Assuming $\alpha > 0$, we will show that $R(A_i) > R(A_j)$. The distribution of A_i is the same as the distribution of $A_j + \alpha$; that is, $F_i(u) = F_j(u - \alpha) \forall u \in \mathfrak{R}$.

$$R(A_i) + c = \int_0^\infty u dF_i(u) = \int_0^\infty u dF_j(u - \alpha)$$

$$= \int_{-\alpha}^\infty (\alpha + v) dF_j(v) \text{ (by change of variable } v = u - \alpha)$$

$$= \int_{-\alpha}^0 (\alpha + v) dF_j(v) + \int_0^\infty (\alpha + v) dF_j(v)$$

$$> \int_0^\infty (\alpha + v) dF_j(v) > \int_0^\infty v dF_j(v) = R(A_j) + c$$

Q.E.D