Title

Evidence for Dynamic Consideration Set Construction in Open-Ended Problems

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors

Nelle, Jonas

Cushman, Fiery

Publication Date

2022

Evidence for Dynamic Consideration Set Construction in Open-Ended Problems

Jonas Nelle (jonasnelle@college.harvard.edu)
Department of Psychology, 33 Kirkland St.
Cambridge, MA 02138 USA

Fiery Cushman (cushman@fas.harvard.edu)
Department of Psychology, 33 Kirkland St.
Cambridge, MA 02138 USA

Abstract

Many decision problems can be divided into three parts: Generating a set of options to consider, evaluating them, and choosing the best. Prior models often assume that the "consideration set" is established in a single step prior to evaluation. Alternatively, people may dynamically and continually assess whether to expand the consideration set based on the quality of the actions considered so far. We use modeling to derive a signature property of dynamic consideration set construction and then demonstrate it in two experiments on human participants.

Keywords: consideration; decision making; option generation

Introduction

In "open ended" decisions, such as what to eat for dinner, we cannot feasibly consider every possible action (Smaldino & Richerson, 2012). Rather, we evaluate a subset of the available actions. How is this "consideration set" constructed?

Much prior work models consideration set construction as a single, pre-deliberative event (Gettys et al., 1987; Morris et al., 2020; Hauser, 2014; Hauser & Wernerfelt, 1990; Kaiser et al., 2013). First a set of options is defined, next each is evaluated, and finally the best is chosen. It is a common experience, however, to have evaluated some options, found none satisfying, and thus generated some *more* for evaluation. The advantages of this "dynamic" strategy are clear: It makes full use of all available information prior to each expansion of the consideration set, rather than pre-committing to the set that seems most advantageous *ex ante*.

A recent study by Callaway et al. (2022) provides suggestive evidence of dynamic consideration set construction. Participants explored a graphical representation of a decision tree by using their mouse to reveal the payoffs at various nodes before settling on a final sequence of actions to execute. The process of node exploration was best described by an optimal model of information search. This model includes both dynamic evaluation of which nodes to search, and also dynamic evaluation of when to terminate search. Our work is similar in spirit, focusing specifically on termination. It extends prior work by investigating spontaneous and fully internalized cognitive search, rather than an externalized analog.

First, following prior work (Sezener et al., 2019; Morris et al., 2020; Callaway et al., 2022), we build a model of dynamic consideration set construction. Next, we use this model to demonstrate a key empirical signature of dynamic consideration set construction: The final action considered tends to be especially high in value, as compared to prior actions that were considered but rejected. This signature is quite intuitive, since the algorithm generates options just until a sufficiently good one is identified. Next, we show empirically that people exhibit this signature during decision-making. We conclude by considering the further steps necessary to develop a unified model decision making: one that fully integrates the processes of option generation and evaluation.

Model

We take as our starting point a recent model of value-based consideration set construction (Morris et al., 2020). An agent must choose between a large but finite set of actions. She has available a noisy representation of each action's value. For each, she can derive a precise representation of its value, but this takes time. The more she deliberates the better able she is to choose the best option. But, deliberation carries an opportunity cost: Time spent deliberating delays the next choice and its reward. The key question is how she can balance these factors in order to maximize expected reward.

Prior work assumed a "static" model of consideration set construction: An agent defines her full consideration set prior to deliberation, next evaluates all actions in that set, and finally chooses the best (Kaiser et al., 2013; Smaldino & Richerson, 2012; Morris et al., 2020). We contrast this with a "dynamic" model in which she can deliberate over each item introduced into the consideration before deciding whether to make an immediate choice, or whether instead to expand the consideration set. Our simulation has two goals. First, we explore when, and to what extent, the dynamic model achieves its greatest advantage over the static model. Second, we identify unique behavioral signatures of each model that guide our subsequent experiments.

The code for these simulations is available at `https://github.com/jonasalexander/thesis`. The agent must choose one of N different possible actions $\{A_1, \ldots, A_N\}$. Each is associated with an approximate representation of its true value $\hat{V}_i = V_i + \varepsilon$, where V_i is the true value and ε is Gaussian noise. (One might interpret this as a model-free or context-free value representation, which is often thought to be an easily retrieved but imprecise estimate). σ^2 is the variance of the noise, $\varepsilon \sim N(0, \sigma^2)$. τ^2 is the variance across actions, i.e. $\hat{V} \sim N(\mu, \tau^2)$.

The agent can enter actions into her consideration set CS

stochastically, in order of decreasing $\hat{V}_i$. This process is instantaneous. But, every action entered into the consideration set must be evaluated, which has some cost $c \geq 0$ (opportunity cost, metabolic cost, etc.). Upon evaluation, the agent learns the precise value of the action V_i. At any time the agent may choose to perform an action in her consideration set, in which case she earns value V_i associated with the chosen action A_i.

Unlike other models, in which V_i is estimated by accumulating noisy unbiased samples (e.g. Callaway & Griffiths (2019)), in our case the agent learns the exact value of the action V_i in one step. Following Morris et al. (2020), we do this for simplicity, but the key insights of our model are easily extended to an accumulator architecture where repeated evaluation of each option is possible.

Static Agent Static agents first generate and then evaluate a fixed number of options. Specifically, they construct a set of the K actions with the highest $\hat{V}$, evaluate all of them at cost $K \times c$, and choose the best. This class of agents is identical to those described in Morris et al. (2020) and also includes take-the-first heuristic agents (Johnson & Raab, 2003) as a special case for $K = 1$.

Dynamic Agent Alternatively, the dynamic agents adjust the size of their consideration set based on the options already evaluated. Following prior work (Callaway et al., 2022), we construe information search as a Markov Decision Problem (MDP) and approximate its solution by Monte Carlo methods. We do not, however, consider this a psychologically plausible model of human decision-making. Rather, we use it to derive empirical signatures of static versus dynamic mechanisms under idealized assumptions.

Because the dynamic agent cannot make a choice without constructing a consideration set of at least one item, the general problem she faces is whether to choose the current best item in her consideration $V_i^b = \text{argmax}_i V(A_i \in CS)$, or whether instead to consider more options.

To start in the simplest case, an agent who has evaluated all options except for the last one ($N-1$ options evaluated) faces the following choice:

$$\begin{cases} V_{N-1}^b \\ \max(V_{N-1}^b, V_N) - c \end{cases} \tag{1}$$

In words, the agent can either choose V_{N-1}^b (thereby terminating search), or can evaluate the next option and potentially find a higher-value one (but will never choose a worse option, hence $\max(V_{N-1}^b, V_N)$). Thus, if the agent has evaluated $N-1$ options, she will evaluate the last option iff

$$E[\max(V_{N-1}^b, V_N) | \hat{V}_N] > V_{N-1}^b + c$$

The expected value of $\max(V_{N-1}^b, V_N)$ is conditioned on $\hat{V}_N$ because in our model this value is known to the agent. The expectation term can be partitioned into the case where the agent falls back on a previous better action and the case where the last action ends up having the highest utility:

$$\mathbb{E}[\max(V_{N-1}^b, V_N) | \hat{V}_N]$$

$$= \int_{x=-\infty}^{\infty} \max(V_{N-1}^b, x) \Pr[V_N = x | \hat{V}_N] \, \mathrm{d}x$$

$$= \Pr[V_N \leq V_{N-1}^b | \hat{V}_N] \cdot V_{N-1}^b + \int_{x=V_{N-1}^b}^{\infty} \Pr[V_N = x | \hat{V}_N] \cdot x \, \mathrm{d}x \tag{2}$$

If we recurse back the the i^{th} decision, we have the following choice:

$$\begin{cases} V_i^b \\ \max(V_i^b - c, \max_{i<j\leq N} V_j - (j-i) \cdot c) \end{cases} \tag{3}$$

where the second term is V^e:

$$V_i^e = \max\left(V_i^b - c, \max_{i<j\leq N} V_j - (j-i) \cdot c\right)$$

Calculating V^e, the expected value of continuing to evaluate more actions, is difficult. There is no simple analytical solution for V^e because each of the actions have different values of $\hat{V}$. In fact, this is the most mathematically significant difference between our approach and sequential search/optimal stopping ("secretary") problems with recall, which assume a constant mean for all options (Ferguson, n.d.; MacQueen & Miller, 1960). Similar to these formalizations, we model this as an MDP where the actions correspond to cognitive operations the agent can perform and the states are defined by the value of the best action so far (V^b) and of the next option ($\hat{V}_{i+1}$) (Callaway et al., 2022). Unlike similar models that myopically compare termination to evaluating just one more option (Gabaix & Laibson, 2005), our model looks ahead to arbitrary depth.

Results

Results are averaged across different parameter settings in order to ensure robust conclusions, as in Morris et al. (2020). Parameters were chosen to be ecologically plausible, to illustrate the key characteristics of the static and dynamic agents, and to be comparable to existing literature. We let N vary from 10 to 100 actions $N = [10, 20, 50, 100]$. The cost of evaluation c ranged from from 0.1 to 10 $c = [0.1, 0.2, 0.5, 1, 5, 10]$, which leads to consideration set sizes in line with empirical estimates of 1-20 (Hauser & Wernerfelt, 1990; Johnson & Raab, 2003). The standard deviation between actions τ varies from 0.5 to 3 ($\tau = [0.5, 1, 3]$), and the level of noise σ applied to $\hat{V}$ varies from 0.5 to 100 ($\sigma = [0.5, 1, 3, 5, 7, 10, 50, 100]$). We repeated this setup for 100 trials per parameter combination, and the agent took 1000 samples to approximate the empirical distribution of $\max(V^b, V)$ for not yet evaluated V and current best value, V^b.

Dynamic Dominance The first key result of our simulations is that the dynamic agent achieves meaningfully better outcomes than the static agent. The size of this advantage depends on the parameter settings. We focus on two especially important factors: σ (the noise applied to $\hat{V}$) and c (the cost of evaluation).

The dynamic agent's advantage is greatest in situations where cognitive search is inefficient (figure 1), i.e. σ is large (relative to τ). Intuitively, when an agent can reliably retrieve the best action in a single step (i.e. $\sigma \ll \tau$ and in the extreme case $\sigma = 0$), the optimal consideration set has size $K = 1$ for both the dynamic and static model. In this case, the dynamic agent cannot achieve better performance than the static agent. As the process of value-guided consideration becomes noisier, there is increasing advantage in terminating search flexibly based on the quality of the options considered so far. However the dynamic advantage starts to decrease again at the other extreme: When σ grows too large relative to τ, consideration becomes random. There is no more information in $\hat{V}$ about V, which degrades the ability of the dynamic agent to accurately predict values of as-yet unconsidered options. Though the advantage decreases, the dynamic agent still performs better because it is able to flexibly terminate search based on the values of actions already evaluated.

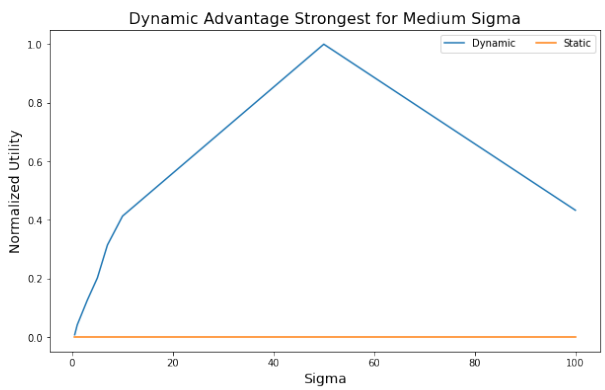

Figure 1: The dynamic advantage as a function of σ. To facilitate comparison, values are normalized so that expected value of static evaluation is always 0 (for all σ), and the maximum expected value across settings of σ is set to 1. At $\sigma = 0$, the static agent and the dynamic agent achieve equivalent performance (both evaluate 1 option, which is guaranteed to be the best). The dynamic advantage grows with σ until, at extremely large values of σ, the relative advantage decreases again.

The dynamic model also shows the greatest relative advantage for intermediate action costs (figure 2). At the extreme of $c = 0$, the optimal agent evaluates all options exhaustively, i.e. $K = N$. On the other hand, when the cost is very high, very small consideration sets are strongly favored because the cost of evaluation outweighs the potential benefits of potentially finding a better option. At the extreme as $c \to \infty$, the optimal agent has $K = 1$. In between, for intermediate action costs, consideration sets of intermediate size are advantaged and the dynamic agent's strengths are most rewarded.

Signatures of Dynamic and Static Processes In order to empirically distinguish static from dynamic consideration set construction, we need distinctive signatures of each process. We demonstrate that a robust signature is the relative value of considered options as a function of their serial position in the consideration process.

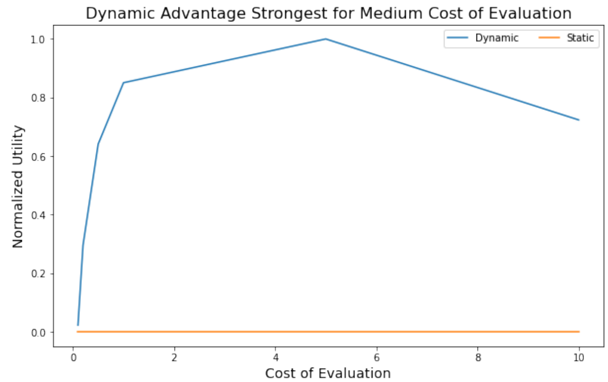

Figure 2: The dynamic advantage as a function of c. To facilitate comparison, values are normalized so that expected value of static evaluation is always 0 (for all σ), and the maximum expected value across settings of σ is set to 1. As $c \to 0$, the static and dynamic agent achieve optimal performance by exhaustively evaluating all options. As c grows larger, both agents evaluate less options and the returns to being able to dynamically construct consideration sets grows. At extremely large values of c, the relative advantage decreases again because it is rarely favorable to evaluate more than one option.

Both classes of agents evaluate options in the descending order of $\hat{V}$, prioritizing early consideration of the most promising candidates. One might expect that the average value of considered options would monotonically decrease with serial order. This is indeed the case for the static agent. Thus, for the static agent, the K^{th} option (the last one evaluated) will tend to have lower value than preceding options.

This is not, however, the case for dynamic agents. Rather, the last option the dynamic agent evaluates tends to be higher in value than the previous options, *because* the dynamic agent is more likely to halt the decision process after having considered a particularly good option. Intuitively, this captures the essential property of the dynamic agent: if it considers a particularly good option, it can terminate the search process. Conversely, if the options evaluated so far are disappointing, the dynamic agent can keep searching.

The dynamic model thus yields unique predictions when comparing it to the static model (figure 3). We exploit this in our experiments on human participants.

Experiments

We conducted two experiments in order to determine whether people dynamically construct consideration sets sensitive to the value of the candidate actions they have considered so far. The first of these experiments involved a reanalysis of data collected for and published in a prior study (Morris et al., 2020). The second provided a conceptual replication and extension of these results. Unlike previous experiments (Sezener et al., 2019; Callaway et al., 2022), our experiments did not externalize the cognitive search process but rather

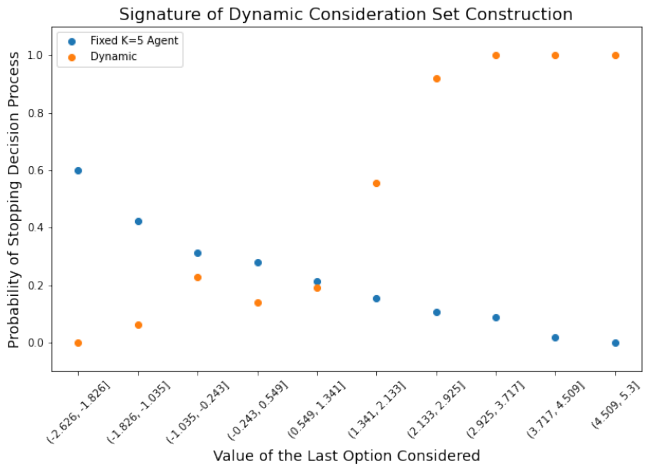

Figure 3: When simulating the considerations sets generated by the dynamic model over a range of parameter settings, we find that, on average, the probability of terminating evaluation at a given an option increases with its value. For a fixed $k = 5$ agent the probability of terminating evaluation at a given option decreases with its value. This is because the agent tends to consider better items first. Here we adopt the exact parameter settings from Morris et al. (2020), performing 10 simulations per setting.

sought to measure the natural, internal consideration set construction processes.

Experiment 1

In the relevant part of Morris et al. (2020), $N = 811$ participants were recruited on Amazon Mechanical Turk and given 25 seconds to name the month of the year the third letter of which, when spelled in English, is closest to "z". After providing their answer, participants were asked to list every month that came to mind—i.e., to report their consideration set. They were asked to list these months in the order that they evaluated them.

While not truly open-ended, limiting the set of total options to the 12 months allowed us to conduct much more rigorous analyses of participants' decision processes. Moreover, we know that few participants evaluated all options (1.7% for Experiment 1 and 10.7% for Experiment 2), so the problem of consideration set construction does in fact arise. Finally, we explicitly set up the task so that no answer was perfect, so participants would always plausibly believe that they could improve their expected point total by investing in further cognitive search.

14 percent of participants in this experiment considered months in a strictly chronological order ("January", "February", "March", etc.), presumably due to familiarity and ease of retrieval. We excluded these participants from the analysis presented here (and we also adopted all the original exclusions of Morris et al. (2020) for other comprehension reasons), but the qualitative results are unchanged, and remain statistically significant, even if we retain them.

We performed two complementary statistical tests on the data. First, we conducted a logistic regression predicting whether or not, following the recall and evaluation of a given word, participants continued cognitive search or instead terminated it (presumably by choosing the best word yet evaluated). The model's predictor was the value of the most recently recalled word. We also included order as a fixed effect in the logistic model in order to rule out its effects as separate from the value itself. A limitation of this approach, however, is that it treats a successive series of decisions by a participant as independent events. Contrary to this assumption, of course, there can be only one termination decision for any given sequence of considered options. To account for this feature of the data we also fit a Cox Proportional Hazards model. For both tests, we found statistically significant results that confirmed our hypothesis (Logistic regression: $\beta = 0.08 \pm 0.01$, $p < 0.001$; Cox proportional hazards: $HR = 1.07[1.05, 1.08]$, $p < 0.001$).

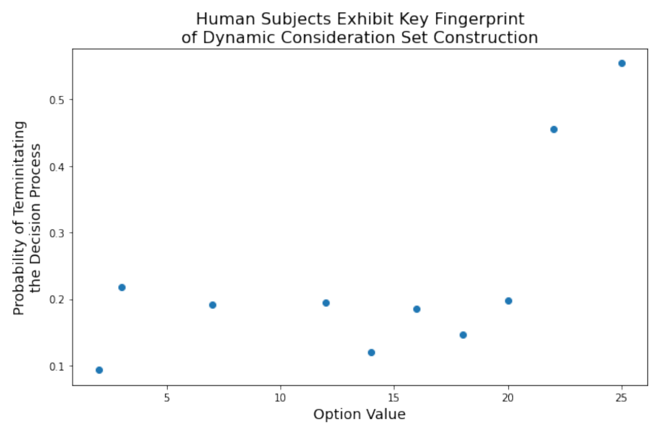

Figure 4: Probability of terminating evaluation at a given option, as a function of that option's value, for Experiment 1.

Although our analysis suggests that participants used a dynamic process to construct and update their consideration set, this experiment has some notable limitations. First, the tendency to consider months in chronological order is an idiosyncratic feature of the experiment orthogonal to the prediction of the static and dynamic models. Even though we removed all participants that considered options in a strictly chronological order, many more considered words in a partly chronological order. For instance, 31 percent of participants considered January first, more than the 8 percent one would expect at random. Second, the distribution of value (i.e., distances of third letters from z) is quite irregular. Our second experiment was designed to avoid these concerns.

Experiment 2

Experiment 2 was modeled on Experiment 1. Participants had to memorize a list of words (this time, a set of semantically unrelated words, in contrast to the 12 months used in Experiment 1). Then, they were asked to name the word (from the

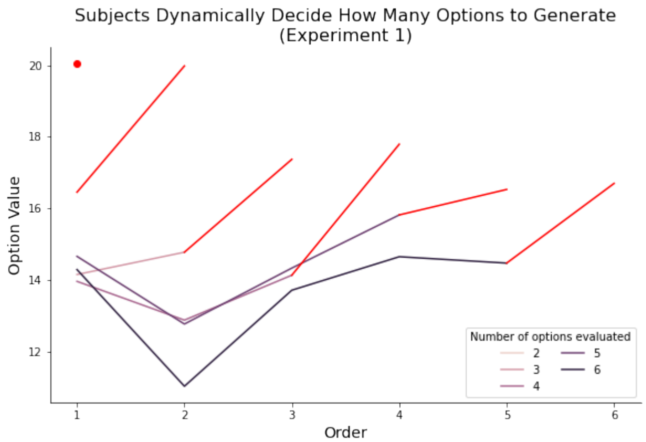

Figure 5: For Experiment 1, the average value of considered items, grouped by the total number of options evaluated, as a function of evaluation order. Blue lines link non-terminal evaluations, while a red line precedes the terminal evaluation. Sequences of evaluation tend to terminate when especially high options are discovered.

memorized list) whose third letter came closest to z (or, in a counterbalanced condition, closest to a). As we intended, participants were not be able to bias retrieval towards higher-value words—i.e., $\sigma \gg 0$.

Methods We presented participants with 13 nouns between 6 and 7 characters long. To encourage memory of the list, we asked them to write a series of sentences incorporating all 13 words. Pretesting of this procedure in a pilot study suggested high (80%) and approximately evenly distributed recall across the 13 words.

Next, we asked participants to name a single word from the list and told them they would receive a monetary bonus depending on how close the third letter of that word was to either "a" or "z", counterbalanced between participants. We constructed the list with the intention of producing a uniform distribution of third letter positions (i.e., one word's third letter was "a", another's was "c", then, "e", "g", etc. through "y"). Due to experimenter error, however, one list had two words with the letter "m" and none with the letter "o".

Finally, we asked participants for the words they had considered in answering the previous question, in the order they had considered them.

We determined our sample size by generating synthetic datasets from the generative models described above: the dynamic model, and a version of the static model in which the consideration set size was determined randomly. To generate these synthetic data, the dynamic model was fit with the following parameters: μ and τ, the mean and variance on $\hat{V}$ were set to be as close as possible to uniform U(1, 26) which is what participants would encounter in the experiment. This meant setting $\mu = 13$ for symmetry and $\tau = 8$ for maximizing the KS-Test metric. We set $\sigma = 10$ for large variance between context-free and context-specific value. Finally, by setting the cost of evaluation to 1, we achieved an average consideration set size of 3-4, in line with empirical estimates (Morris et al., 2020; Hauser & Wernerfelt, 1990). We then performed our intended statistical analyses on these datasets, computing the sensitivity and specificity of those tests when used to infer the underlying data-generating model. When running these simulations with a sample size of 100 at the conventional significance threshold $\alpha = .05$, we had a sensitivity of 100% and a specificity of 98%. To be conservative, we selected a sample size of 200 in our experiment.

Just as in the analysis of Experiment 1, we conducted both a logistic regression and a Cox Proportional Hazards model. In both cases, our independent variable was the value of the most recent word considered and the dependent variable was whether the participant would terminate their cognitive search process.

We excluded participants who failed a simple comprehension question checking their understanding of the reward structure of the task (excluding 30 out of 225, 13%), those whose final answer was not in the consideration set (25 out of 195, 13%), and those who are able to recall fewer than 50% of the words (3 out of 170, 2%). We also excluded from analysis any words present in the consideration set that were not on the memorized list (19 out of 1246, 2%). We corrected participants' spelling errors when and only when they used the correct spelling in the memory check (25 out of 1219 words, 2%); otherwise, these words were also excluded (113 out of 1359, 8%).

All of the above was preregistered on aspredicted.org, with number 78848.

Results As intended, the value of words did not predict whether they were considered; the correlation between value and the probability of consideration was 0.33 ($t(11) = 0.78, p = 0.45$).

Both logistic regression and the Cox Proportional Hazards test indicated that participants were significantly more likely to terminate search following the evaluation of high value words (logistic regression: $\beta = 0.06 \pm 0.02$, $p < 0.001$; Cox proportional hazards: $HR = 1.06[1.04, 1.09]$, $p < 0.001$). This result held when adding order as fixed effect in the logistic regression $\beta = 0.07 \pm 0.02$, $p < 0.001$. (Order is already inherently incorporated in the Cox Proportional Hazards model).

Discussion

Our results suggest a dynamic process of consideration set construction in open-ended problems: People generate a candidate solution, evaluate it, and then decide whether or not to generate another. Thus, in our experiments, the higher in value a considered option was, the more likely participants were to terminate cognitive search.

Although we present a model of this process, it is computationally demanding and thus a poor candidate as an algorithmic-level description of human psychology. Instead,

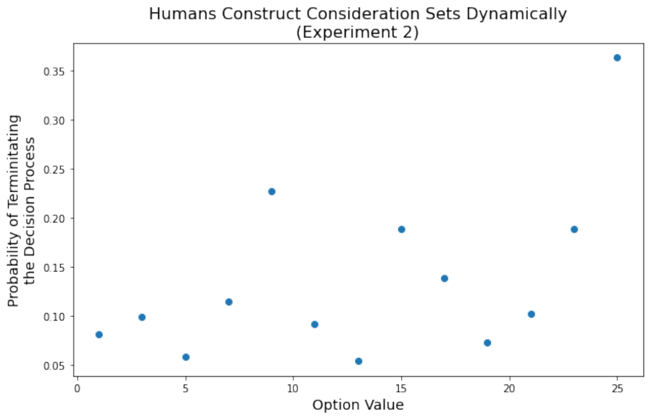

Figure 6: Probability of terminating evaluation at a given option, as a function of that option's value, for Experiment 2.

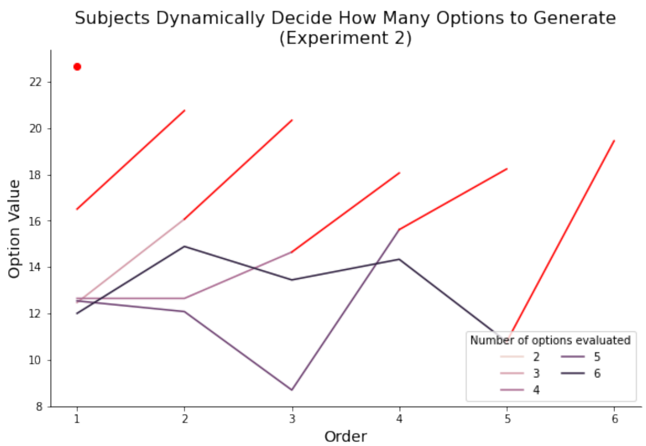

Figure 7: For Experiment 2, the average value of considered items, grouped by the total number of options evaluated, as a function of evaluation order. Blue lines link non-terminal evaluations, while a red line precedes the terminal evaluation. Sequences of evaluation tend to terminate when especially high options are discovered.

we use it to identify an empirical signature of dynamic versus static consideration set construction that we exploit in our experiments.

A key direction for future research, then, is to characterize the computationally efficient algorithms that people use to decide whether to continue or cut off further evaluation (Callaway et al., 2022). We have shown that deriving the optimal stopping point involves a computationally expensive search over a deep decision tree. One family of heuristic approaches would be to approximate the value of continued search via a shallow search—for instance, estimating the expected value of considering just one more item before deciding (Russell & Wefald, 1991; Gabaix & Laibson, 2005). Such a strategy is in keeping with prior research showing that one can often make good decisions about whether to take an action by sampling as few as one possible outcome from a generative model (Vul et al., 2014).

Another direction for future research is to consider different forms of option generation. Here, we focused principally on a value-based approach: i.e., preferentially retrieving those options which have been most valuable in previous contexts (Morris et al., 2020). Another key contributor to the option generation process is semantic knowledge (Zhang et al., 2021). For instance, when attempting to think of what to eat for dinner, one might search semantic associates of "restaurant", "groceries" etc.

Indeed, prior research shows that cognitive search for good actions has a hierarchical structure (Kalis et al., 2013; Klein & Wolf, 1998). For instance, in choosing where to have dinner, one might begin by considering a variety of Chinese restaurants, then a variety of Italian restaurants, etc. We know that hierarchy and abstraction is often important for making decision-making computationally efficient (Russell & Norvig, 2010, p. 406), so this may be an important way in which efficient dynamic consideration set construction can be accomplished. An interesting question, analogous to the one we pursue here, is how people know when to stop "foraging" for options in one semantic space, or at one level of abstraction (Kalis et al., 2013; Klein & Wolf, 1998), and instead switch to another.

Our results also offer an interesting counterpoint to some previous studies of heuristic decision-making. For instance, it has been claimed that a good heuristic approach to decision making is to "take-the-first" option that comes to mind. Simulations performed by Morris et al. (2020) show that this holds under some circumstances, but not under most. As evidence favoring a take-the-first model, Johnson & Raab (2003) and Musculus et al. (2019) show that, when people are given a decision problem, the quality of options generated by those who consider just one option tends to exceed the average quality of options generated by those who consider several. We show, however, that precisely this pattern is predicted by the dynamic model: Those who happen to think of a very good first option cease generating new options, while those who do not continue searching. In other words, deciding to generate only one option does not cause it to have high value; rather, having high value causes a first option to be the only one considered.

The more we learn, the more impressive our decision-making abilities seem. We find that humans dynamically adjust the process of option generation in order to balance the benefits of finding better candidate options against the opportunity cost of delaying action. This illustrates another way in which the study of decision-making must account for the process by which options are generated.

Acknowledgments

We thank Adam Morris for helpful discussion and for generously making available the data used in Experiment 1. We thank Adam Bear and members of the Moral Psychology Research Lab for feedback. We thank three anonymous reviewers for helpful feedback as well. This work was supported by grant N00014-19-1-2025 to FC by the Office of Naval Research.

References

Callaway, F., & Griffiths, T. L. (2019). Attention in value-based choice as optimal sequential sampling. *Preprint*. doi: https://doi.org/10.31234/osf.io/57v6k

Callaway, F., van Opheusden, B., Gul, S., Das, P., Krueger, P. M., Lieder, F., & Griffiths, T. L. (2022). Rational use of cognitive resources in human planning. *Nature Human Behavior*. doi: https://doi.org/10.1038/s41562-022-01332-8

Ferguson, T. S. (n.d.). *Optimal stopping and applications.*

Gabaix, X., & Laibson, D. (2005). Bounded rationality and directed cognition. *Working Paper*. Retrieved from https://scholar.harvard.edu/xgabaix/publications/bounded-rationality-and-directed-cognition-working-paper

Gettys, C. F., Pliske, R. M., Manning, C., & Casey, J. T. (1987). An evaluation of human act generation performance. *Organizational Behavior and Human Decision Processes*, *39*(1), 23-51. doi: 10.1016/0749-5978(87)90044-6

Hauser, J. R. (2014). Consideration-set heuristics. *Journal of Business Research*, *67*(8), 1688-1699. doi: 10.1016/j.jbusres.2014.02

Hauser, J. R., & Wernerfelt, B. (1990). An evaluation cost model of consideration sets. *Journal of Consumer Research*, *16*(4), 393–408. Retrieved from http://www.jstor.org/stable/2489451

Johnson, J. G., & Raab, M. (2003). Take the first: Option-generation and resulting choices. *Organizational Behavior and Human Decision Processes*, *91*(2), 215-229. doi: https://doi.org/10.1016/S0749-5978(03)00027-X

Kaiser, S., Simon, J. J., Kalis, A., Schweizer, S., Tobler, P. N., & Mojzisch, A. (2013). The cognitive and neural basis of option generation and subsequent choice. *Cognitive, affective & behavioral neuroscience*, *13*(4), 814–829. doi: 10.3758/s13415-013-0175-5

Kalis, A., Kaiser, S., & Mojzisch, A. (2013). Why we should talk about option generation in decision-making research. *Frontiers in Psychology*, *4*, 555. doi: 10.3389/fpsyg.2013.00555

Klein, G., & Wolf, S. (1998). The role of leverage points in option generation. *IEEE Transactions on Systems, Man, and Cybernetics*, *28*(1), 157-160. doi: 10.1109/5326.661098

MacQueen, J., & Miller, R. G. (1960). Optimal persistence policies. *Operations Research*, *8*(3), 362-380.

Morris, A., Phillips, J., Huang, K., & Cushman, F. (2020). Habits of thought generate candidate actions for choice. *Manuscript*.

Musculus, L., Ruggeri, A., Raab, M., & Lobinger, B. (2019). A developmental perspective on option generation and selection. *Developmental Psychology*, *55*(4), 745-753. doi: 10.1037/dev0000665

Russell, S., & Norvig, P. (2010). *Artificial intelligence: A modern approach* (3rd ed.). Prentice Hall.

Russell, S., & Wefald, E. (1991). Principles of metareasoning. *Artificial intelligence*, *49*(1-3), 361–395.

Sezener, C. E., Dezfouli, A., & Keramati, M. (2019). Optimizing the depth and the direction of prospective planning using information values. *PLoS computational biology*, *15*(3), e1006827.

Smaldino, P. E., & Richerson, P. J. (2012). The origins of options. *Frontiers in neuroscience*, *6*(50). doi: 10.3389/fnins.2012.00050

Vul, E., Goodman, N., Griffiths, T. L., & Tenenbaum, J. B. (2014). One and done? optimal decisions from very few samples. *Cognitive science*, *38*(4), 599–637.

Zhang, Z., Wang, S., Good, M., Hristova, S., Kayser, A. S., & Hsu, M. (2021). Retrieval-constrained valuation: Toward prediction of open-ended decisions. *Proceedings of the National Academy of Sciences*, *118*(20).