Information Search with Depleting and Non-Depleting Resources

Amber Bloomfield (anbloomfi@umd.edu)
J. Isaiah Harbison (isaiah.harbison@gmail.com)
Susan Campbell (susanc@umd.edu)
Petra Bradley (pscheck@umd.edu)
Lelyn Saner (lsaner@casl.umd.edu)
University of Maryland, 7005 52nd Ave
College Park, MD 20742 USA

Abstract

Predictions about information search behavior have been informed by extensive research in food foraging behavior. However, information foraging environments may differ in key ways from food foraging environments, and these differences may impact search behavior. We investigated the effect of patch distribution (depleting or non-depleting) and ability to return to previously searched patches on participants' decision to switch from one patch to another while searching. Whether or not a participant could return after leaving a patch led to fewer samples and fewer relevant items found. Whether or not the patches depleted and whether it was possible to return to a patch influenced stopping rules, indicating that these factors may alter the size of the increment applied through the Incremental Rule.

Keywords: information foraging, marginal value theorem, stopping rules, patch distribution

Introduction

How do information foragers decide where to look and when to look elsewhere? Research on information foraging has benefitted from an analogy to food foraging behavior (Constantino & Daw, 2015; Dennis & Taylor, 2006; Hills, Jones, & Todd, 2012; Hutchinson, Wilke, & Todd, 2008; Jacobs & Hackenberg, 1996; Kalff, Hills, & Wiener, 2010; Sandstrom, 1994; Wilke, Hutchinson, Todd, & Czienskowski, 2009). The Marginal Value Theorem (MVT), introduced by Charnov (1976), describes how animals foraging for food decide when to cease exploiting a current source to move on to another. In a situation of depleting returns (e.g., a berry bush has fewer berries to offer as a forager consumes berries from that bush), the MVT predicts that a forager will exploit a resource until the point at which the resource's rate of return (i.e., richness) has decreased to the mean return value of the environment. The point at which a forager will leave a resource (i.e., the stopping rule) is then determined by the richness of the resource and the cost of moving on to another resource (the switch cost). A strict interpretation of this rule means that foragers should not leave patches that remain richer than the mean return expected from moving on.

While the analogy to food foraging in animals has proven useful, there are a number of potential differences between the environments in which food and information foraging typically take place. First, in food foraging, a given source normally depletes (e.g., picking a berry results in fewer berries remaining), but depending on the information source, the number of items to search may be so large or replenished so frequently that functionally there is no depletion. For example, the number of tweets on a topic could be so numerous that the information forager is unlikely to read all of them, or the number may increase faster than they could be read as additional tweets are created. Search behavior in environments without depletion has not been investigated in previous studies. A related issue concerns search behavior in environments where it is possible to return to previously investigated patches. While returning to a previous patch is usually sub-optimal in food foraging, as that patch was only left because its richness had diminished beneath the expected return for the environment, a previously visited source may offer the highest richness in the environment if the source does not deplete. Search behavior in environments where it is possible to revisit a patch has not been investigated in previous research.

Stopping Rules

Because the MVT stopping rule requires that a forager have access to information about the richness of a patch at each moment as well as knowledge of the average richness for the environment, more computationally feasible rules have been proposed. The rule which has received the largest share of support in previous information search studies is the incremental rule. In applying this rule, the information forager begins with a set threshold of sampling attempts for a patch (e.g., 10 samples before moving on). When a relevant item is found, the threshold is increased by some increment. People tend to use the incremental rule both for internal search (e.g., memory search; Harbison, Dougherty, Davelaar, & Fayyad, 2009) and external search (e.g., searching in a library or on an online database; Wilke et al., 2009).

Application of the Incremental Rule can produce different switching behavior depending on the size of the increment the information forager uses. For instance, the total failures rule is a special case of the Incremental Rule where the threshold is incremented up by one search attempt for every relevant item found; in essence, samples that yield relevant items do not count against the sampling threshold. An information forager applying the total failures rule will switch patches when the total number of samples not yielding relevant items (i.e., the number of failed samples) exceeds their threshold. Another special case of the Incremental Rule is the total samples rule. In this rule, the size of the increment is zero, such that an information forager will switch patches after a threshold number of samples, no matter how many relevant items are found. Although the total failures rule and the total samples rule are the most common variations of the Incremental Rule, others are possible. For instance, an information forager applying the rule with an increment of size 2 would stay longer in richer patches than a searcher applying either the total failures rule or the total samples rule because each encountered relevant item would increment the sample threshold up by two. These three cases of the Incremental Rule make increasingly similar predictions the poorer the search patch and make the same prediction for patches with zero target items (the case were there would not be an increment).

Predictions

Figure 1 shows the predicted relationship between return rate and number of relevant items found under different increment sizes for the two depletion conditions. These simulated results, like the experiment described below, include three levels of richness: 10%, 30%, and 50% relevant items.

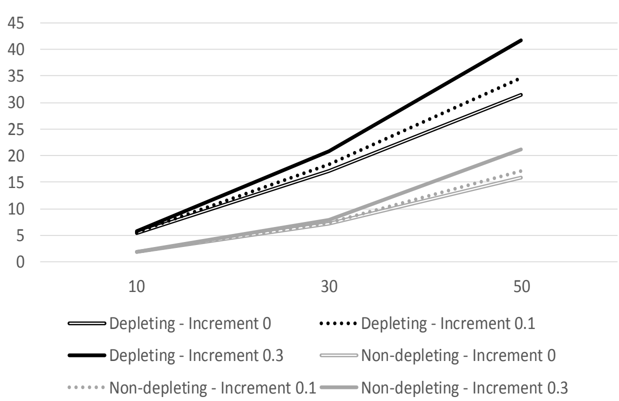

Figure 1: Predicted relationship between return rate and number of relevant items found by depletion condition and increment size.

If participants in our experiment apply the total samples variant of the Incremental Rule (i.e., the size of the increment is zero, represented by the gray and black double lines in Figure 1), they should find more relevant items as patch richness increases, regardless of whether the patches deplete. This is also true for the two other specific cases of the Incremental rule (increments of size 0.1 and 0.3). Therefore, the number of relevant items found does not make a clean distinction between different increment values. However, because of differences in the distribution of relevant items, it is predicted that participants in the Depleting condition will find more relevant items at each level of patch richness than participants in the Non-depleting condition (in Figure 1, the group of black lines are consistently higher than the group of gray lines).

The present experiment also allows us to measure the influence of irrelevant items on search termination in a way that has not previously been explored in the literature. The number of irrelevant items is not accessible to researchers when participants are retrieving from memory because the irrelevant items do not usually elicit a response, and the number of times such irrelevant items are output such that they can be observed is relatively rare (but see Unsworth et al., 2011). In the information search literature, paradigms such as anagram, word search, and fishing provide only a measurement of the time spent not finding relevant words or fish. With the current task, the number of irrelevant items found and their impact on search termination can be measured directly.

The predicted relationship between patch richness and number of irrelevant items found for each depletion condition with three different increment sizes is shown in Figure 2. In this case, there is expected to be a different form of relationship between number of irrelevant items found and return rate dependent on the increment the participant uses. A negative relationship between patch richness and the number of irrelevant items found is predicted if a participant applies the total samples rule (increment of 0). More time is spent sampling relevant items in the richer patches, and since the stopping threshold is not influenced by the number of relevant items found, these relevant items displace the irrelevant items that would be found in poorer patches. Given a larger increment, the relationship should become increasingly positive. Thus, if an information forager applies the total failures rule (increment of 1), a larger number of irrelevant items will be found as the return rate of the patch increases.

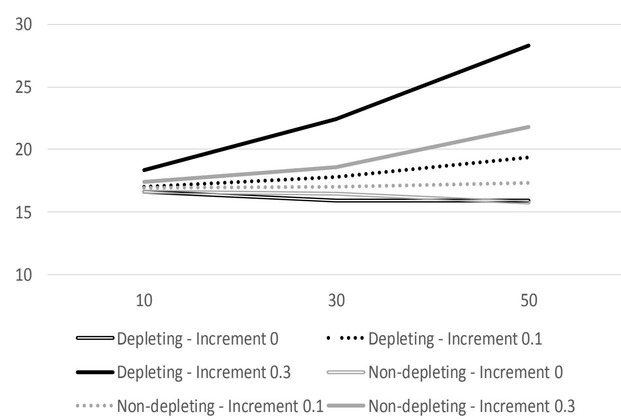

Figure 2: Predicted relationship between return rate and number of irrelevant items found by depletion condition and increment size.

In summary, a different pattern of irrelevant items found in patches of different richness levels will emerge depending on the size of the increment the information forager uses. The predictions are the same for both depleting and non-depleting patches. However, it is possible that information search under conditions where resources do not deplete will lead participants to apply a different size of increment, and so show a different pattern in the relationship between irrelevant items found and patch richness. It may also be the case that whether or not it is possible to return to a patch will impact the increment used to determine when to stop exploiting a given patch. As noted above, previous research has not explored stopping rules when it is possible to return. One possibility is that information foragers will adopt a more conservative approach to stopping when they know they can return to a patch, potentially entailing a smaller increment size. The analyses will explore these possibilities.

Experiment

Search Task

The participants' search task was to collect information about clients for a matchmaking service wherein the participant would use the information found to judge the compatibility of a pair of clients. On each trial, participants saw the name of a client. During the search task, three different sets of tweets (i.e., brief messages stating a piece of information about a person) were offered for the participant to search. Each patch (of tweets) was described as having been produced by a different algorithm designed to gather tweets about the service's clients. Tweets could be about the client (relevant) or about another person (irrelevant).

The richness of the algorithm patches (i.e., the percentage of relevant tweets), could be 10%, 30%, or 50%. The richness levels available among the three patches in a trial varied across trials, with the restriction that all three patches could not have the same return rate. Thus, the highest level of richness available varied from trial to trial. In the depleting condition, the tweets in each of the three patches were ordered such that the rate of return was highest for the first few selections, but decreased as the number of selections increased; in the non-depleting condition, the rate of return remained constant throughout the patch. Patches in the depleting condition had the highest return rate in the first ten selections, and the rate of return decreased in successive sets of ten selections, with the average richness being 10%, 30%, or 50% if the participant sampled 100 times; patches in the non-depleting condition had a constant return rate of 10%, 30%, or 50% throughout.

For each tweet, participants were asked to indicate whether it was relevant for the target candidate or not relevant by clicking one of two buttons. After determining whether a given tweet was relevant to the target, the participant could either choose to view another tweet from the same algorithm or move to a different patch. The switch cost (i.e., the time required to switch to another patch) was 1 second. Depending on condition, participants were either able to return to a patch they had previously searched (return condition), and all three patches remained available throughout each trial, or a patch disappeared after the participant chose to leave the patch (no return condition).

Relevant tweets correctly identified as relevant by the participant were reflected in a running total at the top of the screen, labeled "Score"; participants were not penalized for incorrect decisions, though relevant tweets deemed irrelevant or irrelevant tweets deemed relevant did not increase the participant's score. For each target, participants had two minutes to search for information. At the end of each search, they moved on to the search task for a new target with a new set of patches.

Matching Task

After 20 search trials, each involving a different target person, participants completed a matching task. This task involved 10 trials; on each trial, the participants evaluated the compatibility of a pair of targets (matched based on the gender of match each target was seeking) using the relevant tweets found during the search task. Only relevant tweets the participant had judged to be relevant were available in this phase of the task. Participants indicated their compatibility judgment using an unlabeled slider (scored as a 0-100 scale) with "Horrible Match" and "Great Match" at the poles. The matching task was used to motivate the search task but data from this task were not analyzed.

Participants

Eighty-five participants were recruited via Amazon Mechanical Turk. "Workers" on Mechanical Turk receive monetary compensation for completing online tasks. Workers who selected this experiment task were informed that they would receive $10 for their participation in a task requiring about one hour of time.

Eight participants' data were removed from the sample due to low judgment accuracy, no switching decisions, performing the task twice, or not completing the task. The total analyzed sample, therefore, was 77 participants.

Procedure

Before beginning the experiment, participants were provided instructions and a practice session that included two search tasks and a single matching task. Participants were able to contact the researchers via email if they had questions or encountered difficulty during the experiment. At the end of the experiment, participants were thanked for their participation and paid via Amazon Mechanical Turk.

Participants were randomly assigned to one of the two search conditions: return and non-return. Participants were also, independently, randomly assigned to one of two patch distribution conditions: depleting or non-depleting. There were 16-20 participants in each of the four possible conditions.

Results

As shown in Table 1, the number of samples participants made varied considerably by patch, as did the number of relevant and irrelevant items found.

Table 1. Summary of visits aggregated over patches visited.

	Mean	SD	Min	Max
Total Visits	1.31	1.03	1	10
Total Samples	23.61	16.34	1	100
Total Relevant	9.81	9.37	0	49
Total Irrelevant	13.67	9.80	0	90

Most participants in the Can Return condition chose to return to a previously exploited patch at least once across the ten search trials (~88% of the 41 participants). As shown in Table 2, people were more likely to return in the Depleting condition than the Non-Depleting condition (return was not possible in the No Return condition).

Table 2. Mean visits per visited patch by condition. Standard deviation in parentheses.

	Can Return	No Return
Depleting	1.65 (1.42)	1.00 (0.00)
Non-depleting	1.47 (1.27)	1.00 (0.00)

The analyses in this section focused solely on the participant's first visit (for participants in the Can Return condition) to the first patch offered in each search set, to investigate how depletion condition and return condition influenced participants' search behavior. This focus on the first patch ensured that participants had ample time to investigate the patch prior to choosing to leave, rather than "stopping" as the result of running out of time in the trial. In addition, each analysis involved a multi-level model which estimated an intercept for each participant to account for any individual variability in search behavior.

The first model fit to the data predicted total samples based on return rate of the patch, depletion and return condition. Total number of samples from a patch was related to whether the participant could return (B = 6.98, $p < 0.05$), and the return rate (B = 7.41, $p < 0.05$ comparing 30% to 10% and B = 18.56, $p < 0.05$ comparing 50% to 10%), but not depletion condition (B = -1.98). The means for each condition are shown in Table 3.

Table 3. Mean samples per visited patch by condition. Standard deviation in parentheses.

	Can Return	No Return
Depleting	24.60 (20.89)	28.12 (15.92)
Non-depleting	19.77 (19.1)	28.97 (19.88)

Number of relevant items found was related to all three of the independent variables: return rate (B = 7.04, $p < 0.05$ comparing 30% to 10% and B = 18.20, $p < 0.05$ comparing 50% to 10%), depletion condition (B = -4.84, $p < 0.05$), and return condition (B = 2.60, $p < 0.05$). Average number of relevant items found in each condition are shown in Table 4.

Table 4. Mean relevant items found per visited patch by condition. Standard deviation in parentheses.

	Can Return	No Return
Depleting	13.54 (11.10)	14.72 (10.79)
Non-depleting	8.13 (9.51)	10.13 (10.16)

Only the return condition was significantly related to the number of irrelevant items found by participants (B = 4.16, $p < 0.05$). Participants in the No Return condition found more irrelevant items than those in the Can Return condition (Table 5).

Table 5. Mean irrelevant items found per visited patch by condition. Standard deviation in parentheses.

	Can Return	No Return
Depleting	12.39 (12.92)	13.57 (7.45)
Non-depleting	13.36 (11.34)	19.31 (11.73)

The relationship between number of relevant items found and depletion condition is expected given differences in the distribution of items in these conditions. Because depleting patches were front-loaded with relevant items and their rate of return diminished with repeated sampling, participants in the depleting condition found more relevant items, on average, than did the participants in the non-depleting condition. The significant relationship between return condition and number of relevant and irrelevant items found is more surprising. Participants consistently found fewer relevant and irrelevant items when they were able to return to a previously exploited patch. Taken together, the results suggest that deciding when to switch to a new patch or a previously visited patch, and performing a switch, cost the participants in the Can Return condition a significant number of relevant items and, because there was no penalty for finding irrelevant items, saved them nothing.

Switching behavior

In any given search trial, participants could either switch from the first patch in the set or continue exploiting that patch for the entirety of the two-minute search time. This "sticking and staying" behavior occurred on 9.8% of search trials. We explored how the return rate of the patch and the depletion and return conditions impacted whether or not participants stayed with a patch without exploring any other options in the search trial. Participants were more likely to stay with the first patch when the return rate of the patch was higher (B = -2.57, $p < 0.001$ comparing 30% to 10% and B = -4.89, $p < 0.001$ comparing 50% to 10%), but neither depletion nor return condition impacted likelihood of staying with the same patch (B = -1.36 and B = 0.51, respectively). Participants may have gotten a sense of the highest available return rate over several trials and stayed with the first patch when they suspected this patch had the highest return rate possible.

Stopping rules

The impact of depletion and return conditions on average relevant and irrelevant items found suggests that the Can Return may have led to a less optimal search strategy than the No Return condition. However, the pattern of number of relevant and irrelevant items found across patches of different return rates can provide insight into the dominant increment used in the stopping rule across participants, and whether the increment changed due to return or depletion condition.

Figure 1 shows how the relationship between return rate and number of relevant items found changes depending on return and depletion condition. Consistent with the overall results regarding number of relevant items found, participants in the Depleting condition found more relevant items across all return rates and those in the No Return condition tended to find more relevant items regardless of return rate. Because all forms of the incremental rule predict a position relationship between return rate and the number of relevant items found, these results are consistent with the use of any size of increment.

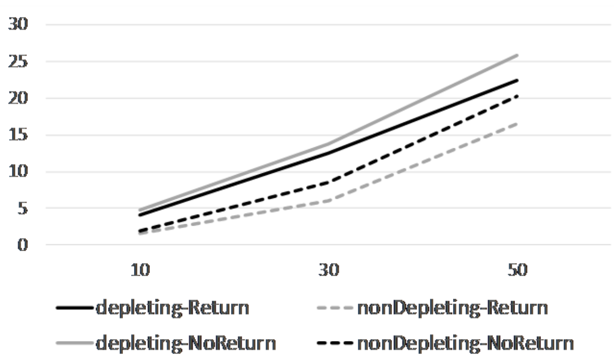

Figure 2. Average relevant items found for each return rate, by depletion and return condition.

Figure 4 shows the average number of irrelevant items found across the three return rates for each depletion and return condition. Across all return rates, participants in the Non-depleting – No Return condition found the most irrelevant items, and there is a slight but positive trend for number of irrelevant items found as return rate increases for this condition and for the Non-depleting – No Return condition. Comparatively, the slope of the Depleting – No Return and Depleting – Can Return appear slightly negative. This difference between the depletion conditions is intriguing because the application of the incremental rule is predicted to result in a different relationship between return rate and number of irrelevant items found depending on the size of the increment the information forager uses.

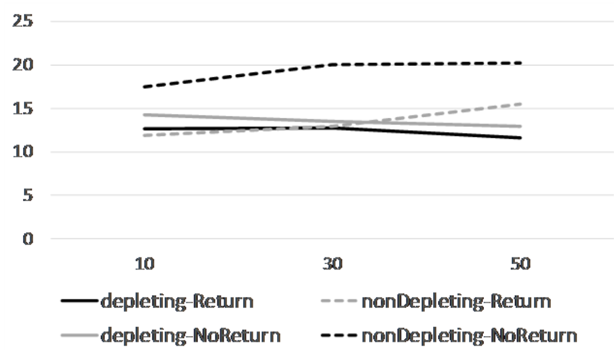

Figure 4. Average irrelevant items found for each return rate, by depletion and return condition.

To explore the dominant stopping rule used by participants in each condition, we fit a linear model predicting number of irrelevant items from patch return rate to each participants' data separately (for all patches, regardless of their position in the search set). Of the 77 participants, there was a non-significant relationship between the number of irrelevant items found and the return rate of the patch for 33. However, 42% of the sample showed a significant negative relationship between number of irrelevant items found and return rate of the patch. Only 14% showed a positive relationship, consistent with using a larger increment. Further, there was variation in the predominant relationship between return rate and irrelevant items found by condition: 63% of participants in the Depletion – No Return condition showed a negative relationship between patch return rate and irrelevant items found; this pattern was found for only 32-40% of participants in the other conditions. Thus, it seems that both whether the patches deplete and whether it is possible to return to previously exploited patches influenced the stopping rule used by participants.

Discussion

The current study explored whether the distribution of relevant items in a patch (depleting or non-depleting) and the ability to return to a previously exploited patch influenced the search behavior of information foragers. Overall, ability to return affected the average number of samples participants made from each patch: those participants who could return made fewer samples, and as a result found fewer relevant and irrelevant items. This suggests that the additional burden of deciding whether to leave a patch to return to a patch previously searched or a patch that has not yet been searched, and the accompanying time taken in switching patches, cut into the amount of time participants in this condition had to exploit any given patch. Because these participants found fewer relevant items compared to participants who could not return, this suggests that performing information search while weighing the ability to return to the current or previously explored patches may lead to suboptimal search. Even in cases where return to a patch is possible, as it is in most real-world information search tasks, it may be advantageous for information foragers to leave each patch only when they are confident they will not wish to return to it.

Across all participants, the number of relevant items found increased with the richness of the patch. This pattern is to be expected if participants are applying the Incremental Rule while searching. Further, the relationship between number of relevant items found and richness of the patch differs only in degree between uses of an incremental stopping rule with different sizes of increments (see Figure 1). However, our study provided an opportunity to directly observe the relationship between number of irrelevant items found and richness of the patch which has not been possible in previous studies. Increments of different sizes are predicted to yield different relationships when the increment is 0, the relationship is predicted to be negative; as the size of increment increases, the relationship is predicted to become increasingly positive (see Figure 2). Of participants who showed a significant relationship between return rate and number of irrelevant items found, the majority showed a negative relationship. This is consistent with application of the total samples rule, which dictates an increment of 0. There were further differences between the conditions: a negative relationship between return rate and irrelevant items found was most prevalent in the Depleting – No Return condition, more common than any other pattern, while other conditions yielded a greater number of positive or non-significant relationships.

Conclusion

Previous research on information foraging has not explored situations where the proportion of targets or relevant items in a patch does not deplete, or cases where it is possible to return to a previously searched patch, yet both scenarios are possible in real-world search tasks. Our results indicate that these factors influence the stopping rules that information foragers apply. Further, our results suggest that having awareness of the ability to return to previously searched patches may lead to less optimal search behavior.

Acknowledgements

This research was supported by the University of Maryland Center for Advanced Study of Language with funding from the Department of Defense.

References

Charnov, E. L. (1976). Optimal foraging, the Marginal Value Theorem. *Theoretical Population Biology, 9*, 129-136.

Constantino, S. M., & Daw, N. D. (2015). Learning the opportunity cost of time in a patch-foraging task. *Cognitive, Affective, & Behavioral Neuroscience*, 1-17.

Dennis, A. R., & Taylor, N. J. (2006). Information foraging on the web: The effects of "acceptable" Internet delays on multi-page information search behavior. *Decision Support Systems, 42*(2), 810-824.

Frank, R. G. (2004). *Behavioral economics and health economics* (No. w10881). National Bureau of Economic Research.

Friedman, M., & Savage, L. J. (1948). The utility analysis of choices involving risk. *The Journal of Political Economy*, 279-304.

Harbison, J. I., Dougherty, M. R., Davelaar, E. J., & Fayyad, B. (2009). On the lawfulness of the decision to terminate memory search. *Cognition, 111*(3), 397-402.

Hills, T. T., Jones, M. N., & Todd, P. M. (2012). Optimal foraging in semantic memory. *Psychological Review, 119*(2), 431.

Hutchinson, J. M., Wilke, A., & Todd, P. M. (2008). Patch leaving in humans: can a generalist adapt its rules to dispersal of items across patches?. *Animal Behaviour, 75*(4), 1331-1349.

Jacobs, E. A., & Hackenberg, T. D. (1996). Humans' choices in situations of time-based diminishing returns: Effects of fixed-interval duration and progressive-interval step size. *Journal of the Experimental Analysis of Behavior, 65*(1), 5-19.

Kalff, C., Hills, T. T., & Wiener, J. M. (2010). Human foraging behavior: A virtual reality investigation on area restricted search in humans. In *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (pp. 168-173).

Parsons, G. R., & Kealy, M. J. (1995). A demand theory for number of trips in a random utility model of recreation. *Journal of Environmental Economics and Management, 29*(3), 357-367.

Sandstrom, P. E. (1994). An optimal foraging approach to information seeking and use. *The Library Quarterly*, 414-449.

Thorelli, H. B., & Engledow, J. L. (1980). Information seekers and information systems: a policy perspective. *The Journal of Marketing,* 9-24.

Unsworth, N., Brewer, G. A., & Spillers, G. J. (2011). Factors that influence search termination decisions in free recall: An examination of response type and confidence. *Acta psychologica, 138*(1), 19-29.

Wilke, A., Hutchinson, J., Todd, P. M., & Czienskowski, U. (2009). Fishing for the right words: Decision rules for human foraging behavior in internal search tasks. *Cognitive Science, 33*(3), 497-529.