Experiments on Information Foraging

David J. Bryant
DRDC – Toronto Research Centre

Defence Research and Development Canada
Scientific Report
DRDC-RDDC-2017-R159
November 2017

IMPORTANT INFORMATIVE STATEMENTS

Endorsement statement: This publication has been peer-reviewed and published by the Editorial Office of Defence Research and Development Canada, an agency of the Department of National Defence of Canada. Inquiries can be sent to: Publications.DRDC-RDDC@drdc-rddc.gc.ca.

This report was completed under the Joint Intelligence Collection and Capability (JICAC) project (JFD 4.1), sponsored by the Chief of Defence Intelligence (CDI).

Abstract

Two experiments tested predictions of Information Foraging Theory (IFT) pertaining to "patch-leaving," the decision to abandon an information source to search elsewhere. IFT predicts that increasing the between-patch cost associated with moving from one source of information to another should lead foragers to increase the time spent in each patch. Similarly, increasing the within-patch cost associated with processing individual information items should likewise increase the time spent in each patch. Participants searched for information relevant to solving a series of simulated analysis questions using the INformation FOraging Cognitive Analysis Tool (INFOCAT) platform. Information items were separated into a number of discrete databases and participants were allowed to freely search information and select items they judged to be relevant. In Experiment 1, the time delay associated with opening a database (between-patch time) was varied across analysis questions, while in Experiment 2, the time delay added to opening an information item (within-patch time) was varied across analysis questions. The results of the two experiments indicated no evidence that participants' patch-leaving decisions were affected by either between- or within-patch costs. There was also no indication that participants' search behaviour changed over the course of an experimental session. These results suggest that people do not necessarily apply optimal foraging strategies to information search. Although further research is needed to conclusively determine whether peoples' information foraging conforms to predictions of IFT, an opportunity exists to enhance the efficiency and effectiveness of the intelligence analysis process by translating IFT concepts into training and decision support for intelligence analysts.

Significance to defence and security

Information gathering is a significant part of the military intelligence analyst's role and efforts to better understand and support this activity could result in great benefits to the analyst's overall performance. Information Foraging Theory offers a way to understand information seeking as an adaptive process. The experiments reported here test whether predictions of this theory are upheld by human information foragers in a simulated analysis task. This line of research could provide valuable insight into the cognitive strategies of analysts and guide development of decision support concepts to enhance their effectiveness. In particular, supporting optimal information foraging can help analysts engage in more efficient search and achieve either a greater yield of information for time spent foraging.

Résumé

Deux expériences ont vérifié les prédictions de la théorie du butinage des renseignements (TBR) concernant la décision d'abandonner les recherches dans une source d'information pour poursuivre les recherches ailleurs. La TBR prédit que l'augmentation du coût associé au déplacement d'une source d'information à une autre devrait inciter les butineurs à passer plus de temps dans chaque source d'information. Dans un même ordre d'idées, l'augmentation du coût associé au traitement d'éléments d'information individuels devrait aussi inciter les butineurs à passer plus de temps dans chaque source d'information. Les participants ont fait des recherches d'information afin de répondre à une série de questions d'analyse simulées à l'aide de l'outil d'analyse cognitive du butinage de renseignements (INFOCAT). Les éléments d'information ont été divisés en un certain nombre de bases de données distinctes. Les participants pouvaient mener leurs recherches librement et choisir les éléments qu'ils jugeaient pertinents. Dans l'expérience 1, le délai d'ouverture d'une base de données (pendant la transition d'une source à une autre) variait selon les questions d'analyse, alors que dans l'expérience 2, le délai d'ouverture d'un élément d'information (pendant la recherche dans une source) variait selon les questions d'analyse. Les résultats des deux expériences n'ont donné aucune preuve que le coût de la transition ou de la recherche avait une incidence sur la décision de passer à une nouvelle source d'information. Rien n'indique non plus que les participants ont modifié leur comportement de recherche au cours des deux expériences. Ces résultats portent à croire que les gens n'appliquent pas nécessairement des stratégies de butinage des renseignements optimales. Bien que d'autres recherches soient nécessaires pour établir avec certitude que le butinage de renseignement est conforme aux prédictions de la TBR, il est possible d'améliorer l'efficacité du processus d'analyse du renseignement en convertissant les concepts de la TBR en formation et en aide aux décisions pour les analystes du renseignement.

Importance pour la défense et la sécurité

La collecte d'information est une partie importante du rôle de l'analyste du renseignement militaire, et les efforts pour mieux comprendre et appuyer cette activité pourraient apporter de grands avantages au rendement général de l'analyste. La théorie du butinage des renseignements offre une façon de comprendre la recherche d'information comme un processus adaptatif. Les expériences dont nous rendons compte ici vérifiaient si les prédictions de cette théorie étaient accomplies par des butineurs de renseignements humains dans une tâche d'analyse simulée. Cet axe de recherche pourrait fournir des données précieuses sur les stratégies cognitives des analystes et orienter le développement de concepts d'aide aux décisions pour améliorer leur efficacité. Plus particulièrement, la promotion d'un butinage des renseignements optimal peut aider les analystes à réaliser des recherches plus fructueuses et à obtenir une plus grande quantité de renseignements pour le temps consacré au butinage.

Table of contents

List of figures

List of tables

Acknowledgements

The author would like to thank Tonya Hendriks, Brenda Fraser, Sarah Gibbon, and William Kozey for their invaluable help in conducting this research.

This page intentionally left blank.

1 Introduction

1.1 Background

1.1.1 Intelligence analysis

Military intelligence analysis is a complex, cognitively demanding process. The objective of analysis is to gain information superiority through knowledge of an operational area to support informed decision making [1]. This requires extensive information and so a great deal of an intelligence analyst's time is spent gathering data that will be used to build situation awareness [2]. Information gathering is considered one of the two major loops (the other being sensemaking, the iterative process of interpreting information, generating and testing hypotheses, and determining additional information needs) that make up the intelligence process [3][4].

Information gathering is currently a labour-intensive activity requiring analysts to work under conditions of information overload and severe time pressure [5]. Military analysts generally consider information gathering to take up too great a portion of their time, leaving insufficient time for sensemaking [6][7]. One reason analysts spend so much time gathering information is that they have access to a wide range of information sources. Information overload and time pressure dramatically increase the cognitive burden placed on analysts and leave them more susceptible to cognitive biases and error [7][8].

In response to this concern, research was initiated under the Joint Intelligence Collection and Analysis Capability (JICAC) project to investigate information search from the perspective of Information Foraging Theory (IFT) [9]. IFT describes optimal behaviour in information search tasks, establishing a normative baseline with which to compare the performance of analysts. It also provides guidelines for evaluating cognitive search strategies that can potentially enhance the performance of analysts. Because information search comprises a large part of an analyst's work effort, measures that enhance the efficiency of analysts in this activity should provide a significant benefit to analysts' overall performance. This report describes the theoretical basis of this research and the results of two experiments examining information search in a simulated intelligence analysis task. Further research directions are discussed in conclusion.

1.1.2 Information Foraging Theory (IFT)

Developed in large part through the work of Peter Pirolli and his colleagues, IFT is a framework for explaining how humans search for and exploit information [10][11][12]. In this framework, information search strategies are evaluated with respect to their suitability within the environment in which search takes place [10, p.17][13]. Drawing on the study of optimal foraging behaviour, IFT likens information gathering by humans to the search for food engaged in by organisms. In this analogy, the human is an "informavore" seeking and consuming information [10, p. 13]. Moreover, as with an animal seeking food, a human informavore is expected to attempt to derive the greatest cumulative value in information while expending the least energy/time possible—that is, to optimize his/her foraging performance.

Information foraging is conceived of as a process of "moving" (both physically and in terms of directing attention) from location to location (typically referred to as "patches" in foraging theory), seeking valuable information, periodically exploiting that information, then moving on [11]. The forager gains information but pays costs in terms of the effort and time spent on moving and exploiting. The goal of the forager is to obtain the best ratio of gain to cost possible [10, pp. 30–31].

Information gain can be thought of as the usefulness derived from information [10, p. 21]. Usefulness, however, is not a consistent quality but, rather, depends on the objectives of the individual using the information [2]. To have a more objective measure, information value can be defined with respect to the relevance of information to search goals of the forager [3]. In this context, relevance is measurable by the degree of overlap between some piece of information and the explicit or implicit information goal of the forager. This definition of relevance has the advantage of being objectively measurable in terms of the semantic similarity between texts [14].

The costs of information foraging come in the form of the physical and cognitive effort of seeking information as well as the time spent on that activity [11]. Time is especially important as information foraging imposes what are called opportunity costs; that is, foraging causes one to forfeit benefits that could be gained by engaging in other activities [10, p. 13]. Generally, the physical and cognitive costs of foraging are assumed to be captured by time (i.e., both are correlated with time, which serves as a proxy measure) so as to represent foraging costs in the single unit of time [15].

IFT can illuminate many aspects of human information search, such as how people determine what information is worth exploiting, where to look for information, and when to stop looking in a patch and move to another [9]. The question of how a forager determines when to move on from one source of information to another is the problem of "patch-leaving" [16]. An information environment will be patchy to the extent information items occur in a non-uniform fashion that causes subsets of items to be grouped or clumped, as a result of different physical media or the organization of information by the creators of those media. Optimal foraging depends on making good patch-leaving decisions to ensure that attention is allocated such that the forager has the greatest likelihood of finding useful information. In particular, a forager should leave a patch when the likelihood of finding addition useful information drops below a certain point [17].

Depending on the specific nature of the forager's task and environment, there may be different cognitive strategies to achieve this goal but, generally speaking, optimal foraging is whatever maximizes the ratio of gain to cost. IFT serves as a guide to determining what those cognitive strategies are and whether a given strategy will result in optimal performance under the conditions of the environment [11].

1.2 Patch-leaving

The research described in this report focuses on patch-leaving and the question of whether human information foragers behave in ways that are consistent with predictions of IFT. Foraging costs are typically divided into two main components, the so-called "between-patch" activities of looking for, and movement to, the next place to forage (e.g., accessing a new database, pulling out a folder of documents, etc.), and "within-patch" activities, involving the exploitation of resources within a patch (e.g., viewing items contained in a database, reading a paragraph within a document, etc.) [10][11]. The gain derived from foraging is affected by both these types of costs, which must be considered in determining the optimal time to leave a patch.

This is accomplished with Charnov's Marginal Value Theorem (MVT), which describes the relationship between the instantaneous, or marginal, rate of gain within a patch and the average rate of gain within the environment as a whole [10, pp. 35–36][18]. As illustrated in Figure 1, the cumulative gain from foraging within a patch, $g(t_W)$, is an increasing function of time in which the rate of increase declines over time. This pattern of diminishing returns, in which cumulative gain is always increasing—but at a decreasing rate—is assumed to describe most natural foraging environments (e.g., [17][18]). Beyond this assumption, the precise form of the cumulative gain function will vary according to the specific resource values and processing costs associated with items within a patch. The rate of gain is captured by the slope of a line tangent to any given point on the gain curve. When the gain function provides diminishing returns over time, as in Figure 1, that slope is largest for points at the beginning of the gain curve and decreases steadily over time.

The MVT states that a forager will maximize his/her long-term rate of gain by leaving a patch when his/her current rate of gain equals the long-term average rate of gain across all patches in the environment [18]. The optimal time at which to leave a patch can be found by drawing a line beginning at the average between-patch time, t_B, prior to entering the patch, tangent to the gain function. The point of tangency indicates the optimal leaving time, t^*, assuming that the gain function represents the average of patch gain functions in the environment [10, pp. 35–36]. At that point, the rate of gain within the patch equals the average expected rate of gain to be obtained from search within the environment. Any subsequent foraging within the patch will yield a lower rate of gain than leaving and seeking another patch.

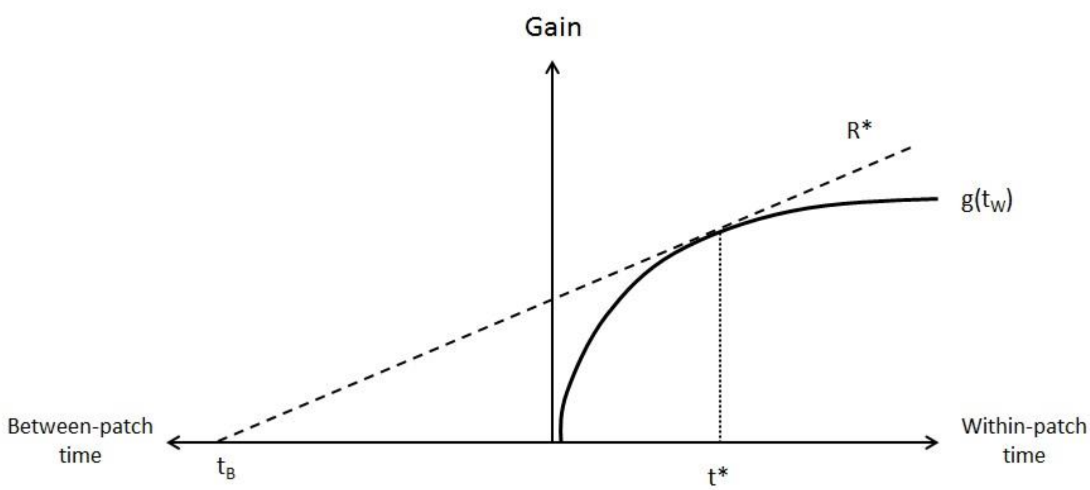

Figure 1: *Illustration of diminishing returns and Charnov's Marginal Value Theorem.*

A number of environmental factors affect the optimal patch-leaving time for a forager. First, the richness of patches—the average rate of gain associated with patches in the environment—determines whether a longer or shorter patch-leaving time is optimal [10, pp. 36–37]. This factor is illustrated in Figure 2, which shows two hypothetical gain functions. The first gain function, $g_1(t_W)$, represents an environment in which patches have a relatively low rate of gain overall. The second gain function, $g_2(t_W)$, represents an environment in which patches, on average, result in higher gains. As shown in Figure 2, assuming the same between-patch time required to move between patches, the second gain function is associated with a shorter optimal leaving time. Thus, when foraging in a relatively rich environment, in which patches on average offer higher gains, a forager is expected to exhibit shorter leaving times than a forager in a relatively poor environment. Note that the magnitude of the gain rate is also affected by the within-patch, or processing, time. Decreasing the average amount of time required to process each resource item has the effect of increasing the magnitude of the gain function. Thus, as average within-patch time increases, a forager is expected to exhibit longer patch-leaving times.

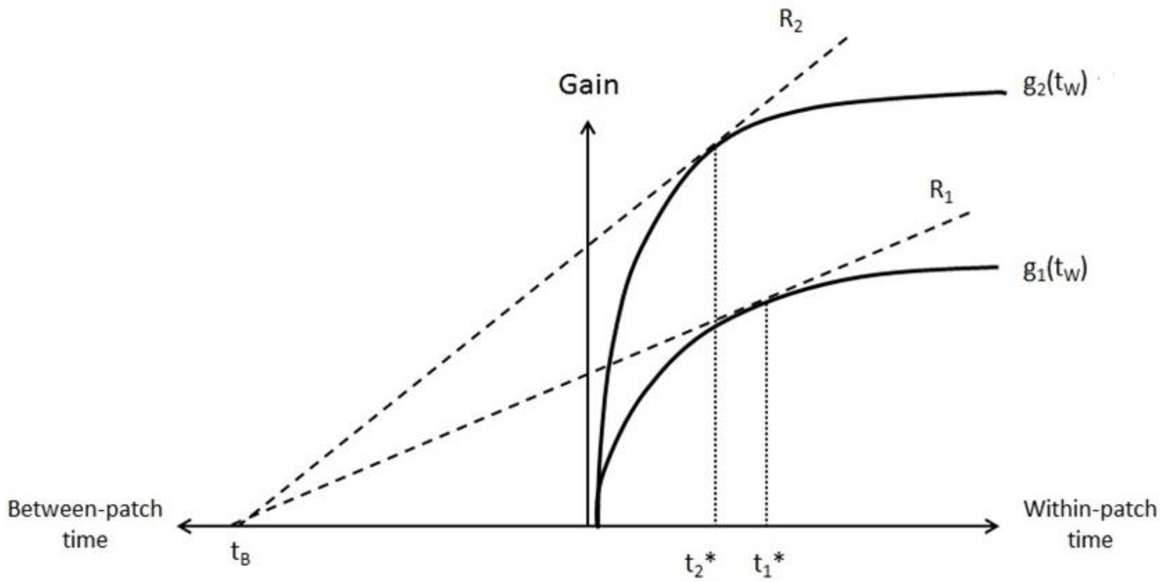

Figure 2: *Effect of patch richness on optimal patch-eaving time.*

Optimal patch-leaving time is also affected by the cost of moving between patches [10, pp. 36–37]. Figure 3 illustrates two environments in which patch richness is represented by a single gain function, $g(t_W)$, but the between-patch time required to move between patches is either relatively long, t_{B1}, or short, t_{B2}. The longer between-patch time yields a shallower tangent to the gain function, resulting in a longer optimal leaving time, whereas the shorter between-patch time yields a shorter optimal leaving time. Thus, as the average time required to move from one patch to another increases, a forager is expected to remain for longer amounts of time in each patch.

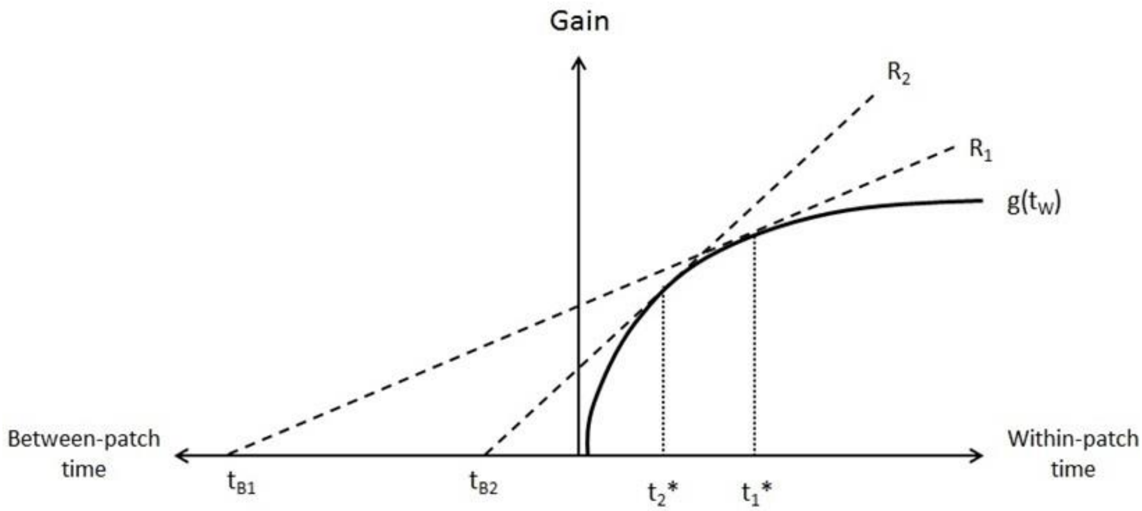

Figure 3: *Effect of between-patch time on optimal patch-leaving time.*

1.3 Information foraging by humans

Given the claims of IFT, the question arises as to whether there is any evidence that human information foragers behave in ways that are consistent with optimal foraging. This question is particularly important for intelligence analysts whose job performance depends on being effective information foragers. Pirolli and others [10][11][19] have argued that the evolution of human intelligence has been, to some degree, driven by the need to develop sophisticated strategies for finding and obtaining food (see also [20]). As a result, one would expect humans to exhibit behaviour consistent with optimal foraging in information search tasks.

IFT provides a framework in which to identify optimal information search behaviour. It identifies key factors associated with costs and benefits that define optimal cognitive strategies. If people are naturally sensitive to these factors, it will be relatively easy to integrate optimal strategies into intelligence analysts' information seeking procedures. The evidence concerning how well IFT predicts human search behaviour, however, is mixed.

Some previous research supports the claim that people behave adaptively when performing information processing tasks. Payne, Duggan, and Neth [16], for example, found people were sensitive to the continuous rate of gain while performing a task and also tended to switch tasks when the rate of gain dropped below some threshold. They related their results to foraging heuristics described by Stephens and Kreb's [21] for food foraging. In another study, Payne et al. [16] found that participants were sensitive to factors such as rate of return and seemed to switch between tasks at times consistent with a heuristic patch-leaving process [22].

On the other hand, some research has suggested humans are not necessarily optimal foragers. Importantly, this research, conducted by Puvathingal [23], directly investigated information foraging by intelligence analysts. In two experiments, trained analysts and novices searched for information that could be used to evaluate potential causes of the 1989 turret explosion onboard the *USS Iowa* battleship. The information was accessible through three databases, in which the participant could scan a series of labelled links to individual information items. When clicked, each link opened to display the full item in a separate screen along with a scale along which the participant could rate the usefulness of the item. Participants were allowed to freely interact with the databases, in any order and for any amount of time, but the overall time allowed for the task was limited. Puvathingal varied the lengths of delays required to open databases (between-patch time) and individual items (within-patch time) to determine whether these factors would affect participants' patch-leaving decisions. As noted above, increasing both the between- and within-patch time is expected to lead to longer patch-leaving times. Results indicated that neither analysts nor novices exhibited leaving behaviour consistent with optimal foraging.

1.4 Purpose

This report describes experiments that investigate whether people's information foraging behaviour conforms to predictions of IFT. The experiments follow on those of Puvathingal [23] and test the hypotheses that longer between- and within-patch costs lead people to remain foraging within patches longer. Despite recent research (e.g., [19][23]), it remains unclear the extent to which human search behaviour conforms to predictions of optimal foraging theory. The experiments described here examine the impact of travel (between-patch) and item processing (handling) costs on peoples' information search behaviours. These are important factors that determine optimal patch-leaving in foraging and are expected to affect peoples' decisions on how long to remain searching for information in a given location. The two experiments were motivated by a need for further evidence in the context of intelligence analysis.

Information search was studied in the context of a simulated intelligence analysis task, using the INformation FOraging Cognitive Analysis Tool (INFOCAT), which was developed to present simulated intelligence analysis problems and collect participant data [24]. Three simulated analysis problems were developed along with information items derived from various sources, including published reports and news articles. The INFOCAT platform provides an organizational structure in which information is separated into a number of discrete locations and allows the participant to search through the organizational structure to locate information useful in addressing the question or problem. For each analysis problem, participants were allowed to freely search information and select items they judge to be relevant to addressing the analysis problem.

2 Method

Two experiments were performed following the same general method.

2.1 Participants

Twenty-four individuals participated in total (12 in each experiment). All participants were active duty members of the CAF who had no experience in any intelligence occupation.

2.2 Design

The two experiments employed 2-within factors mixed designs. In both experiments, the first within-participant factor was the sequence of analysis questions. Participants completed three analysis problems set up with the same distribution of information items (i.e., the proportion of relevant to irrelevant information was constant) across databases. This was intended to allow participants to gain knowledge of the underlying distribution of information items in databases so that changes in information foraging behavior between the initial (no knowledge), second (intermediate knowledge), and final problem (most knowledge) could reveal whether participants internalized knowledge about this aspect of the information environment.

In Experiment 1, the second within-participant factor was the time delay associated with opening a database (between-patch time), which was varied across analysis questions. All databases within an analysis question had the same level of delay, either 4, 8, or 16 seconds added to the opening of a database. The order of database delay times over the three analysis questions was alternated across participants, with half receiving the 4s, 8s, and 16s delays for the first, second, and third questions respectively and the other half receiving the 16s, 8s, and 4s delays for the first, second, and third questions respectively. The delay associated with opening each individual information item (within-patch time) was held constant at 2s for all items in all databases.

In the Experiment 2, the second within-participant factor was the time delay added to opening an information item (within-patch time), which was varied across analysis questions. All items within an analysis question had the same level of delay, either 2, 4, or 8 seconds added to the opening of each information item. The order of item delay times over the three analysis questions was alternated across participants, with half receiving the 2s, 4s, and 8s delays for the first, second, and third questions respectively and the other half receiving the 8s, 4s, and 2s delays for the first, second, and third questions respectively. The delay associated with opening a database (between-patch time) was held constant at 2s for all databases in all analysis questions.

The lengths of the between-patch and within-patch delays were chosen to ensure that participants could not exhaustively search the databases and, thus, examine every item in all databases. Given the 20 minute overall time limit, participants would be confronted with a situation in which they might, at most, be able to consult half of the total number of items. This design feature created time pressure like that experienced by intelligence analysts and forced participants to make strategic decisions concerning how long to spend foraging in each database.

2.3 INFOCAT platform

The experiments were conducted using the INFOCAT platform [24], which was developed to create an experimental task that: a) poses a question or problem to an experiment's participant that requires an information search; b) provides an organizational structure, in which information is separated into a number of discrete locations; and c) provides an organizational structure through which the participant can search to locate information useful in addressing the question or problem.

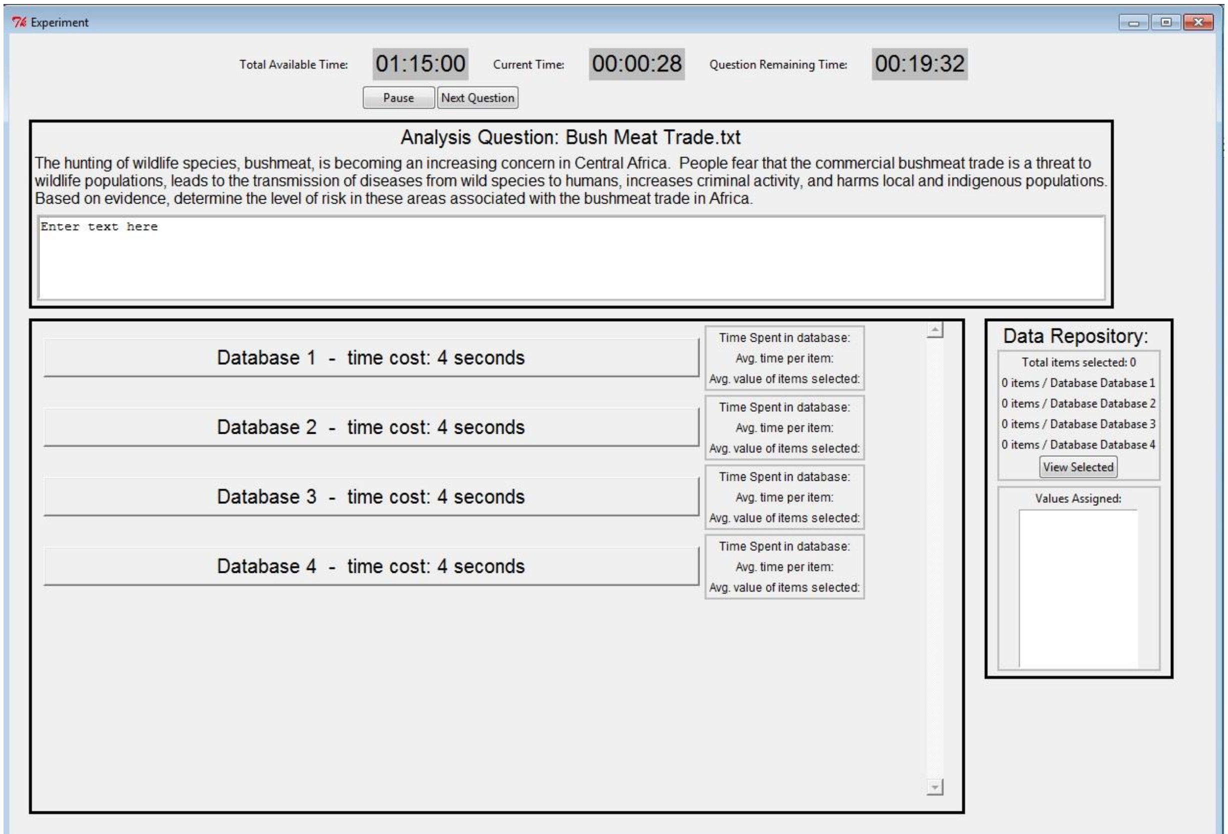

Figure 4: *The INFOCAT main interface.*

The main INFOCAT interface, illustrated in Figure 4, has four main components. The first, at the top of the display, is a set of timers that display the overall experiment time, the time in each analysis question, and the remaining question time, as well as the buttons to pause the experiment and move to the next question. The second is a field, under the timers, in which an analysis question is displayed along with an area in which the participant can input his/her interpretation of, or formal response to, the analysis problem. The second component, located at the lower left, contains fields for individual "databases" that can be accessed to search for information. A database is simply a link to another field in which a subset of information items can be accessed. Thus, this area comprises the organizational structure in which information items are located and each database can be considered an information patch. In the experiments, participants could access four databases by clicking a labelled button on the screen. The third area, to the lower right, is a "data repository," which is a field in which a participant has access to a record of all information items found and selected in searches through databases.

Figure 5 shows the interface that appears when a participant clicks on a database. The section showing links to the databases has been replaced with an area containing links to all the individual information items available in the selected database. In both experiments, the links showed only the item number so participants had no cues as to the content of any item. Participants could read an item by clicking its link, which called up a screen, shown in Figure 6, containing the item and two buttons that participants used to select the item for storage in the data repository or reject the item. Selecting an item flagged it for storage in the data repository area of the interface, which allowed the participant to review all previous selections at any time. After selecting or rejecting an item, a 7-point Likert scale appeared, asking the participant to rate the relevance of the item to the analysis problem.

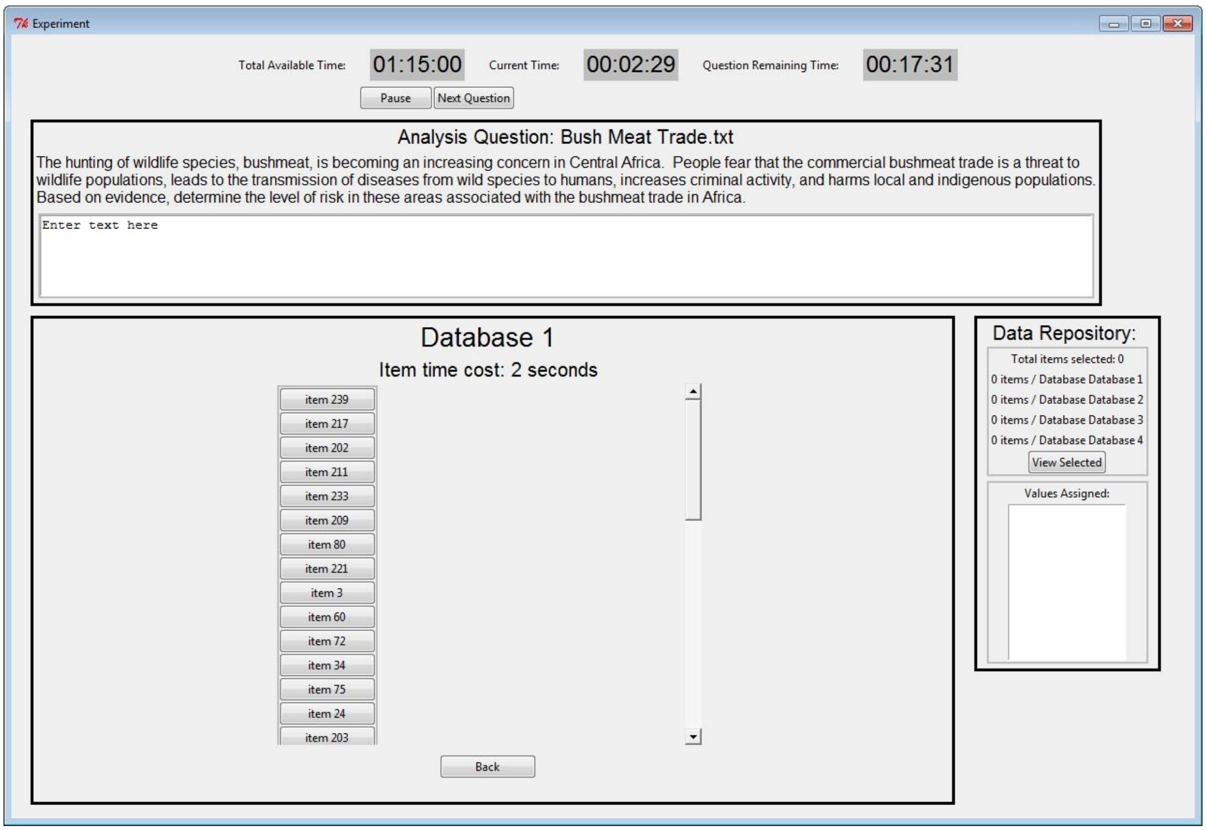

Figure 5: *INFOCAT database interface.*

Information items could be revisited at any time during a session. Participants were able inspect a previously inspected item and make a different judgment (select or reject) and a different relevance rating. The number of views and all associated judgments and ratings were recorded for analysis.

2.4 Procedure

Participants completed the same procedure in both experiments. Upon being briefed about the general purpose of the experiment and signing the consent form, participants were instructed in the use of the INFOCAT platform with an example analysis problem. The experimenter showed participants how to access information items using the point-and-click interface, how to select items for use, and how to indicate ratings of relevance for items.

Each analysis problem was a moderately complex statement of a political/social issue for which an analyst could develop a brief estimate (statement of issues and likely course of events). Three questions were created: a) the cause of a 1989 explosion onboard the *USS. Iowa*, b) the potential risks of migration into Romania, and c) the potential consequences of the central African bushmeat trade.

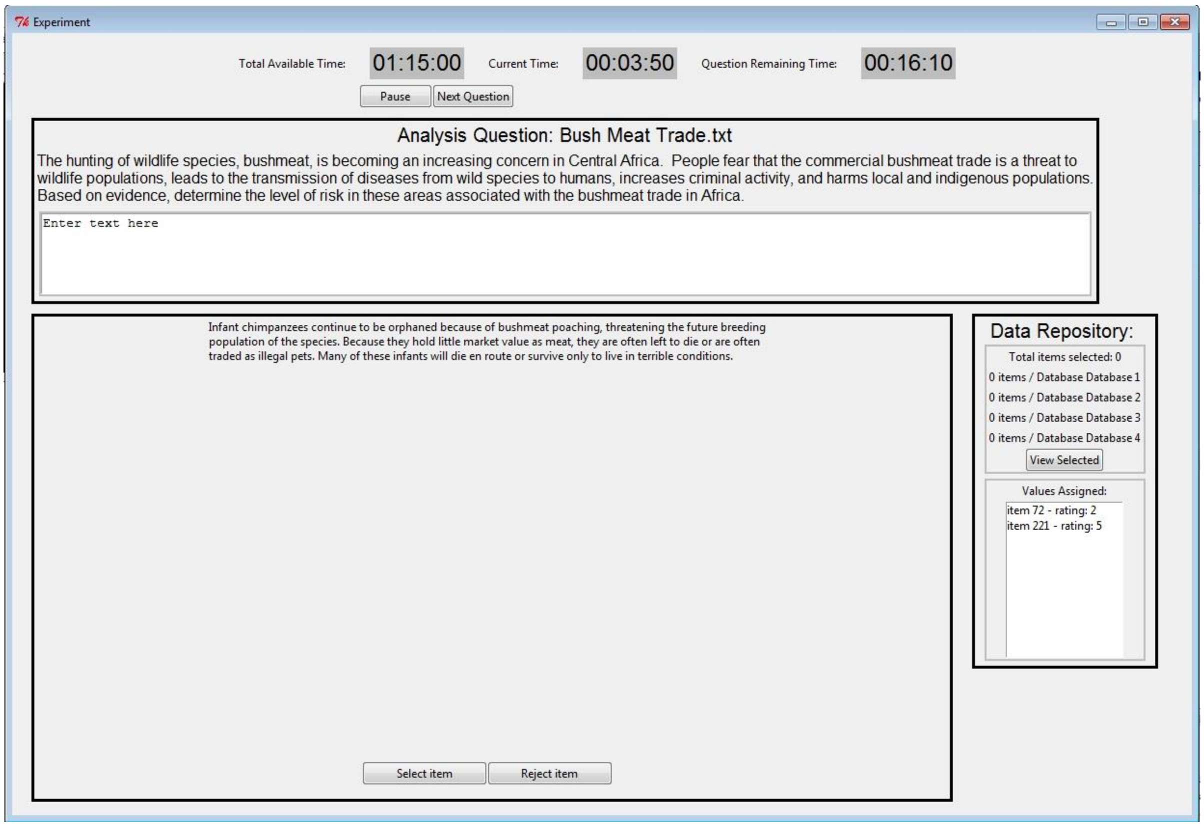

Figure 6: *INFOCAT item interface.*

Each question had an associated corpus of 250 information items, which were short (typically 3–4 sentences) statements conveying some information pertaining to the topic of the analysis problem. These items were rated on a 7-point Likert scale by four independent judges for relevance or usefulness in addressing the corresponding analysis problem. Items were developed such that the majority of items were of low relevance, which received ratings of 1 or 2. Items were considered to be relevant to answering an analysis question if they received an average rating of 4 or higher from the judges.

Information items were organized into four databases containing a total of 40 items each. Items were selected randomly for each participant by the INFOCAT platform from the total available set of items. For all analysis questions, the four databases contained 7, 8, 12, and 13 relevant items (rated 4–7 by judges) with the assignment of the number of relevant items to each database being counter-balanced across participants. All remaining items in a database were irrelevant (those receiving average ratings of less than 3). The numbers of relevant items were chosen to approximate a normal distribution of relevant items across the four databases. Different numbers of relevant items in each database created a strategic problem for participants who would want to spend more time in richer databases (with 12 or 13 relevant items) than poorer ones (with only 7 or 8 relevant items). The INFOCAT interface provided no indication of database richness but participants could estimate the quality of a database on the basis of their pattern of successes and failures in finding relevant items as they foraged a database. The number of relevant items in each database was held constant across analysis questions for each participant (but counter-balanced across participants). Thus, if Database 1 contained 13 relevant items in the first question, then Database 1 would contain 13 relevant items in the second and third questions as well. This consistency was intended to allow a participant to notice the correlation between database and number of relevant items, which could be used to guide foraging decisions in the second and third questions.

For each analysis problem, participants first read the full problem, taking as much time as required to understand it. Participants then began the information search task. Using the INFOCAT interfaces described above, participants searched for information items. They were instructed to select information items that they believed were useful in answering the analysis question and reject those they believed had no use. Participants made ratings of item relevance on a 7-point Likert scale for all items viewed. Participants were free to search however they wished, selecting databases and items within databases in any order. The only restriction on their search was a 20 minute time limit, after which all information items, other than those selected, would become inaccessible. After the 20 minute search session, the participant was asked to write a short brief on the problem. Participants were allowed a 10 minute interval to write their brief, during which they were able to refer to selected information. Writing the briefs maintained the analytic nature of the task for participants but the briefs themselves were not analysed.

3 Predictions

Theoretical predictions of information foraging performance were derived to provide a baseline against which to compare participants' observed behaviour. Because the information items and database structure were the same in the two experiments, the cumulative gain function—the total amount of information value acquired as a function of item selection—is the same in both cases. Value, in this context, corresponds to the judged relevance ratings associated with information items. An item's judged rating was assumed to be a quantitative measure of its value to the participant. The average value of relevant items, V, was computed from the relevant items assigned to databases in all conditions.

The experiments were designed as tests of the hypotheses that people will adjust their time spent in databases based on the database-entry delays according to predictions of Charnov's MVT [18]. For the MVT to be applicable, the cumulative gain function of patches must exhibit a pattern of diminishing returns. It was initially assumed that having a limited number of relevant information items in each database would be sufficient to produce such a cumulative gain function. An analysis after the experiments had been conducted, however, revealed that the actual cumulative gain function for the databases used in the experiments did not exhibit a pattern of diminishing results. Unlike a natural environment, in which a forager typically must spend more time per item as it depletes the resources of a patch, the processing time for items in the databases did not change as items were selected.

The cumulative gain function was determined by calculating the expected value at each item selection, from 1 to 40 in a database. The expected value was equal to the sum of the probabilities of acquiring 1 to 10 (maximum number of relevant items on average in the four databases) relevant items at each item selection point, weighted by the average value of a relevant item. The probabilities of acquiring relevant items were determined by a hypergeometric distribution with 40 total items, a total of 10 relevant items, across 1 to 40 item selections, with 1 to 10 potential "successes." The resulting predicted gain function shown in Figure 7 indicates the amount of information value expected on each selection within a database for which the participant has no knowledge (i.e., no knowledge of the exact number of relevant items, the ordering of items within the database, or the value of any item).

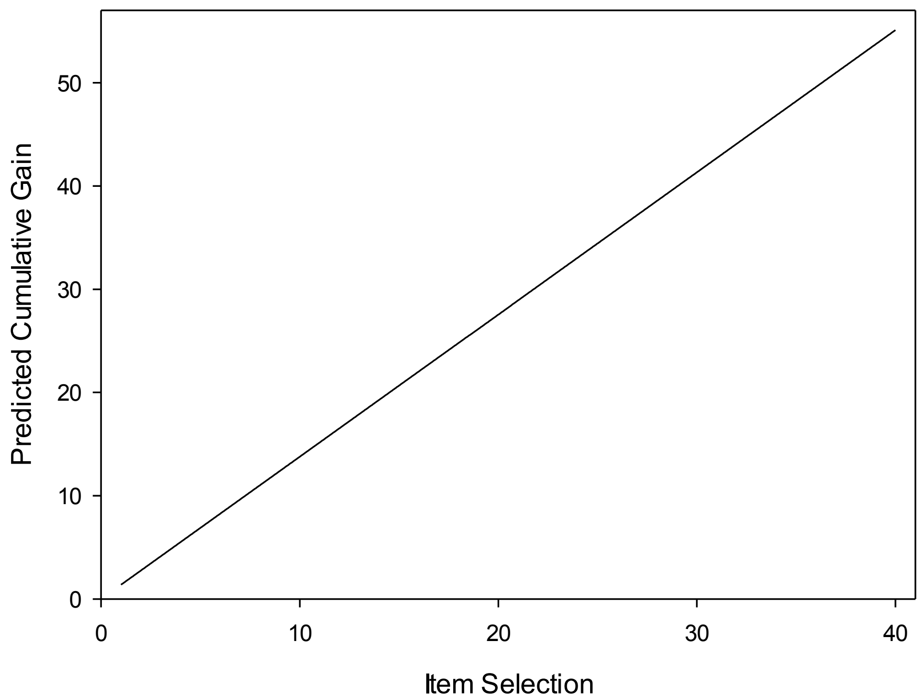

Figure 7: *Cumulative gain function for databases in both experiments.*

As shown in Figure 7, expected cumulative gain is a linear function of item selection for the information structure used in the two experiments. Thus, a participant could expect to acquire the same increment in value with each item selection. Assuming that each selection takes the same amount of time (i.e., participants do not systematically require increasing amounts of time as a function of item selection and all variability in item processing time is due to participant factors rather than structural factors associated with the information environment or INFOCAT interface), this results in a flat, linear function of expected rate of gain (value per unit time) with item selection. Note, this analysis is not affected by the item-opening delay, which is constant for all items within an analysis question.

A linear gain function violates one of the assumptions of Charnov's MVT, namely the condition of diminishing returns in which the rate of gain declines with time spent foraging in a patch [18]. Consequently, the prediction of the MVT—that a forager will depart a patch when the instantaneous rate of gain in the patch falls below the overall environmental rate of gain—cannot occur. Because the theoretical expected rate of gain in a patch is constant, there is never a point at which it will fall below the environment average. Instead, an optimal forager is expected to exhaustively search each database until depletion before moving to the next database.

Although the expected gain rate is a flat function of item selection, the addition of a time delay to access a database distorts that function somewhat. The database-access delay is a cost that reduces the expected gain rate but the impact of that cost diminishes as a forager continues to select items within a database. The expected rate of gain for an item selection was calculated by the following formula:

$$g_i = \frac{\lambda \cdot \mathrm{V}}{T_B + T_W} \qquad (1)$$

Where, g_i is the cumulative gain expected with the ith item selection (out of 40 possible selections in a database), λ is the summed probability of encountering 1 to 10 relevant items with the ith selection as derived from the hypergeometric distribution, V is the average value of a relevant item, T_B is the total between-patch cost (assumed to be the database access delay), and T_W is the total within-item cost (assumed to be the cumulative item processing time over i selections). Thus, as illustrated in Figure 8, the expected rate of gain in both experiments will show an initial steep rate of climb over the first few selections before leveling off.

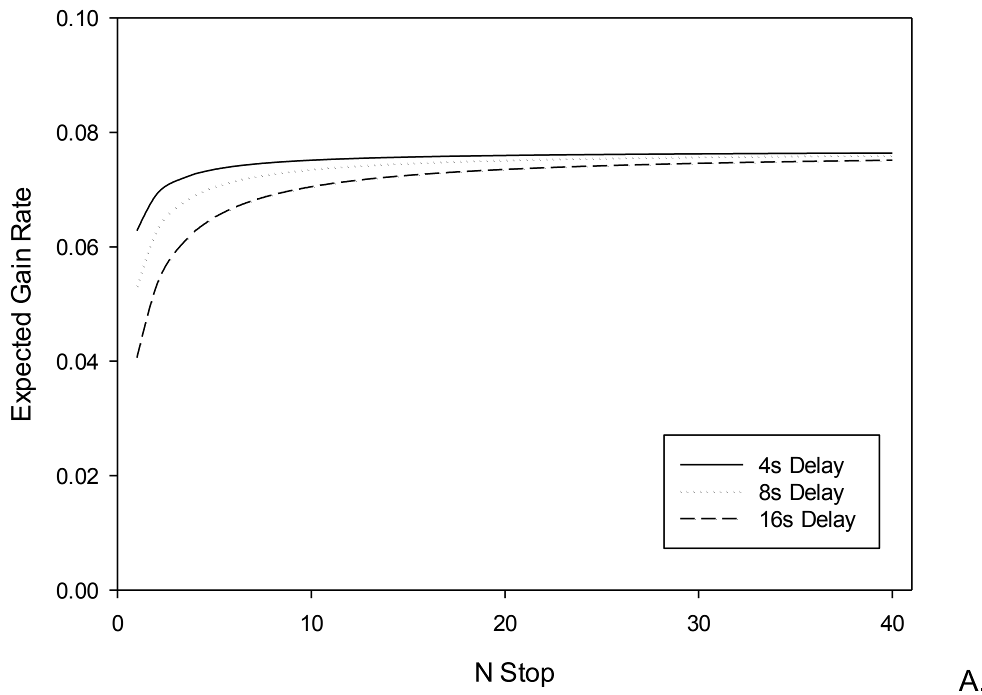

Figure 8: *Predicted rates of gain (value per second) for A) Experiment 1 and B) Experiment 2.*

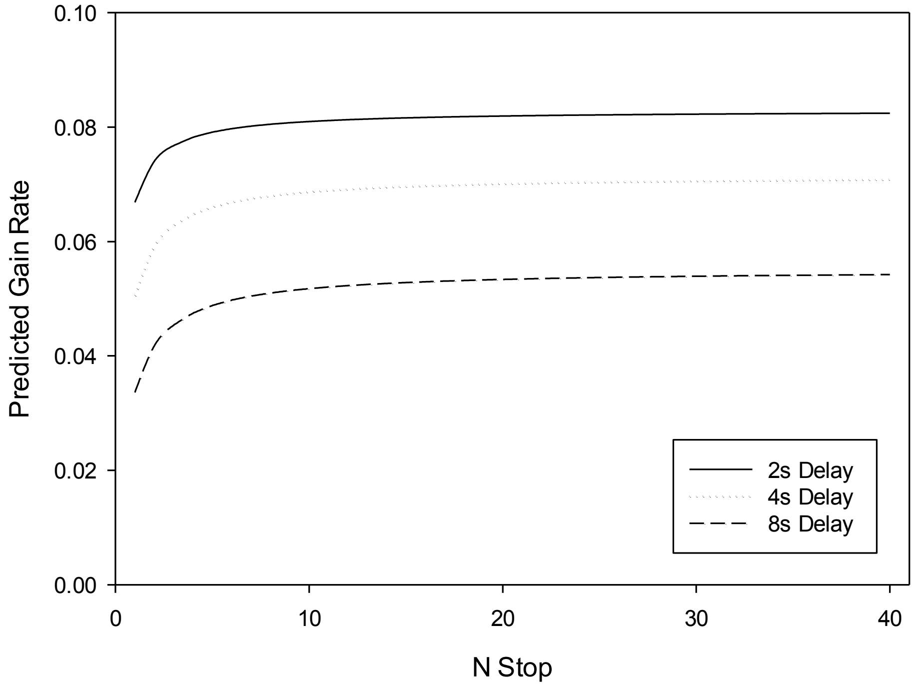

B.

Figure 8: *Predicted rates of gain (value per second) for A) Experiment 1 and B) Experiment 2 (cont.).*

In Experiment 1, the length of the database-access delay is varied and Figure 8A shows that the longer the delay, the more gradual the leveling off. In Experiment 2, the database-access delay was held constant but the delay associated with accessing each item was varied and Figure 8B shows that the shape of the gain rate curve in each item delay condition is the same but the overall gain rate is different.

The gain rate functions indicate that, when database-access cost is considered, leaving a database too soon (within the first 5 or 6 selections) will be highly sub-optimal. Beyond that point, although the theoretical optimal stopping point is the 40th selection, there is little difference in expected outcome between any database-leaving point, assuming the forager can return to any database if foraging time is available or the forager has access to a practically unlimited number of databases of the same type.

This analysis of the information environment used in the two experiments indicates that it does not present a powerful test of the predictions of the MVT. Charnov's MVT assumes that patches exhibit a pattern of diminishing returns such that the rate of gain obtained by foraging within a patch decreases as a function of time [18]. This is not true in the experimental environments where rates of gain were constant.

4 Results

4.1 Experiment 1: Results

4.1.1 General foraging behaviour

A striking feature of these results is the large degree of variability exhibited by participants in all measures of performance. As shown in Table 1, participants opened, on average, between three and four databases in each condition but there was no reliable difference between conditions [$F(2,20) = 0.21$, $MS_e = 99$, *n.s.*]. Participants opened roughly the same total number of information items in all three conditions [$F(2,20) = 0.21$, $MS_e = 117.4$, *n.s.*]. Participants selected (indicated as relevant) slightly more items in the 4s delay than other conditions but this difference was not statistically significant [$F(2,20) = 1.94$, $MS_e = 7.69$, *n.s.*]. Participants spent slightly more time processing individual items in the 8s and 16s than 4s delay conditions but this difference was also not statistically significant [$F(2,20) = 0.93$, $MS_e = 55.36$, *n.s.*]. Finally, mean ratings given to selected items were similar across all conditions [$F(2,20) = 1.11$, $MS_e = 0.01$, *n.s.*].

4.1.2 Patch-leaving

Table 2 shows the mean amount of time spent by participants in a database for the three delay conditions. These mean times do not include the time spent on the delay to access a database. On average, the time spent in a database tended to be higher in the 16s than 4s and 8s delay conditions and only slightly higher in the 4s than 8s delay condition. These differences, however, were not statistically reliable [$F(2,20) = 1.75$, $MS_e = 9978$, *n.s.*].

Participants' tendency to remain in a database was examined by comparing the number of items viewed by participants during a database visit across delay conditions. As shown in Table 2, the mean number of items opened per database visit exhibited the same pattern as database visit times—highest in the 16s delay condition, followed by the 4s delay condition, and lowest in the 8s delay condition—but these differences were not statistically significant [$F(2,20) =0.84$, $MS_e = 41.08$, *n.s.*].

Another way to assess the tendency of participants to remain in foraging patches is the mean giving-up time (GUT) across conditions. The GUT is the amount of time a participant spent in a database after the last success (selecting a relevant item) before leaving the database. As shown in Table 2, participants exhibited longer GUT, on average, in the 16s delay than all other conditions. However, mean GUT was longer in the 4s delay than 8s delay condition, violating predictions of the MVT. Moreover, the variability of GUTs was very large, meaning no differences achieved statistical significance [$F(2,20) = 1.14$, $MS_e = 1389$, *n.s.*].

The giving-up number (GUN)—the number of items opened between the last success (item selected) and the decision to leave the database—was also considered. Table 2 shows that there was essentially no difference in the mean GUN between conditions [$F(2,20) = 0.04$, $MS_e = 4.36$, *n.s.*]. The average GUN of less than three items indicates that, in all conditions, participants tended to not persist in searching within a database for very long after an individual success, although there was considerable variability in participants' giving-up decisions.

Because the GUT and GUN are measured from the last selected item to the point at which the participant leaves the database, they are not strictly dependent on the number of relevant items in the database. The number of relevant items in a database may only be pertinent when the forager hits the overall time limit of the analysis question and must terminate search. Analyses of GUT and GUN with such terminal cases removed confirmed the initial findings.

Overall, there is no indication that participants stayed longer in databases, or had any greater tendency to delay leaving a database, in the longer delay conditions.

Table 1: *Measures of general foraging performance (Experiment 1).*

	Databases Opened		**Items Opened**		**Items Selected**		**Time per Item (s)**		**Selected Item Rating (1–7)**	
Delay	**Mean**	***S.D.***	**Mean**	***S.D.***	**Mean**	***S.D.***	**Mean**	***S.D.***	**Mean**	***S.D.***
4 s	3.75	*1.14*	70.42	*31.51*	22.08	*12.82*	16.41	*8.20*	5.52	*1.08*
8 s	4.00	*1.65*	69.00	*32.34*	18.50	*9.59*	18.66	*9.33*	5.56	*0.64*
16 s	3.25	*1.36*	67.58	*29.50*	18.42	*6.50*	18.74	*10.54*	5.43	*1.06*

S.D. Standard deviation

Table 2: *Measures of patch-leaving behaviour (Experiment 1).*

	Time in Database (s)		**Items Opened per Database**		**Time in Database After Last Selection (s)**		**Items Opened in Database After Last Selection**	
Delay	**Mean**	***S.D.***	**Mean**	***S.D.***	**Mean**	***S.D.***	**Mean**	***S.D.***
4 s	351.21	*114.02*	20.34	*10.87*	51.10	*46.88*	2.52	*1.90*
8 s	330.25	*118.56*	19.64	*11.29*	42.12	*33.18*	2.68	*1.99*
16 s	449.17	*270.41*	22.85	*10.15*	66.11	*65.71*	2.72	*1.71*

S.D. Standard deviation

4.1.3 Analysis question order

The number of relevant items in each database was held constant across analysis questions for each participant (but counter-balanced across participants). Thus, if Database 1 contained 13 relevant items in the first question, then Database 1 would contain 13 relevant items in the second and third questions as well. Potentially, a participant who noticed this correlation between database and number of relevant items could use that information to guide foraging decisions in the second and third questions, possibly altering the participant's observed behaviour.

Table 3 shows data across analysis question order as opposed to delay condition. There is no indication that participants' search behaviour changed over the course of the experimental session as no statistically significant differences were observed across question order. Thus, participants opened roughly same number of databases [$F(2,20) = 0.84$, $MS_e = 1.28$, *n.s.*] and items per database visit [$F(2,20) = 0.58$, $MS_e = 42.92$, *n.s.*] regardless of question order. Time spent in databases also did not differ as a function of question order [$F(2,20) = 0.67$, $MS_e < 0.01$, *n.s.*]. Participants selected slightly fewer items per database visit for the second question (this was always the Romania Migration question) than the first or third but this difference was not significant [$F(2,20) = 0.54$, $MS_e = 12.57$, *n.s.*]. Participants spent roughly the same amount of time processing items [$F(2,20) = 1.67$, $MS_e = 22.55$, *n.s.*] and gave selected items roughly the same ratings, on average, at all question orders [$F(2,20) = 0.09$, $MS_e = 0.04$, *n.s.*].

Neither the GUT [$F(2,20) = 1.58$, $MS_e = 1540.2$, *n.s.*] nor GUN [$F(2,20) = 0.88$, $MS_e = 3.76$, *n.s.*] differed significantly from the first to third question.

4.1.4 Information gain rate at the time of leaving (ATOL)

The gain rate at the time of leaving (ATOL) is the rate of value obtained per time spent achieved by a forager in a database at the point at which the forager exited that database. Gain rate ATOL is often compared to the overall average rate of gain achieved by a forager in an environment as a test of whether the forager adheres to the prediction of the MVT; i.e., an optimal forager should exhibit a gain rate ATOL slightly lower than the environmental average. In this experiment, however, the information environment does not meet the assumption of diminishing returns demanded by the MVT, so this comparison is not appropriate. Instead, the mean gain rates ATOL for the three delay conditions were computed and compared to one another to determine whether participants exhibited any differences in patch-leaving behaviour as a function of database-access delay. Such differences could potentially reflect participants' perceptions of the environmental gain rate for different levels of between-patch cost. As can be seen in Table 4, the observed gain rate ATOL was larger in the 4s than 8s and 16s delay conditions, suggesting that participants tended to leave databases earlier in 4s delay condition. The observed difference, however, was not significant [$F(2,20) = 1.96$, $MS_e = 0.01$, *n.s.*].

Table 3: *Measures of foraging behaviour across analysis question order (Experiment 1).*

	Databases Opened		Items Opened per Database		Items Selected per Database		Time per Item (s)		Selected Item Rating (1–7)	
Question	**Mean**	**S.D.**	**Mean**	***S.D.***	**Mean**	***S.D.***	**Mean**	**S.D.**	**Mean**	***S.D.***
First	3.42	1.16	22.05	*10.36*	7.11	*4.65*	19.23	8.70	5.48	*1.10*
Second	4.00	1.65	19.64	*11.29*	5.70	*4.55*	18.66	9.33	5.57	*0.64*
Third	3.58	1.38	19.48	*9.13*	6.84	*6.34*	15.92	9.99	5.47	*1.04*

S.D. Standard deviation

Table 4: *Observed gain rate ATOL (Experiment 1).*

	Gain Rate At Time Of Leaving (Cumulative Rating per second)	
Delay	**Observed**	***S.D.***
4 s	0.146	*0.152*
8 s	0.079	*0.045*
16 s	0.084	*0.041*

S.D. Standard deviation

4.2 Experiment 2: Results

4.2.1 General foraging behaviour

As in the first experiment, participants exhibited large variability on all measures of performance (see Table 5). Participants opened, on average, more databases than did participants in Experiment 1 but again no reliable difference between conditions was found [$F(2,20) = 1.58$, $MS_e = 1.89$, *n.s.*]. In contrast to the first experiment, participants did show a reliable difference in total number of items opened, with participants opening more items in the 2s than 4s delay condition and more in the 4s than 8s delay condition [$F(2,20) = 12.58$, $MS_e = 193.8$, $p<.05$]. This seems to reflect a limiting effect of the item-opening delays (i.e., the longer this delay, the less time overall a participant would have to open items), as well as differences in the amount of time spent on each individual item, minus the item-opening delay, across conditions. Participants spent considerably more time processing individual items (not including the delay imposed on opening each item) in the 8s than 4s delay condition and more in the 4s than 2s delay condition [$F(2,20) = 6.35$, $MS_e = 34.81$, $p<.05$]. Participants selected more items in the 2s than 4s delay condition and more in the 4s than 8s delay condition but this effect was not statistically significant [$F(2,20) = 3.45$, $MS_e = 53.34$, *n.s.*]. Finally, mean ratings given to selected items were similar across all conditions [$F(2,20) = 0.54$, $MS_e = 0.18$, *n.s.*].

4.2.2 Patch-leaving

Table 6 shows the mean amount of time spent by participants in a database for the three delay conditions. In contrast to the first experiment, participants showed fairly large differences in the amount of time spent per database, but not in a clear pattern. Participants spent more time per database in the 2s than 8s, and more in the 8s than 4s, delay conditions. This effect was nearly significant [$F(2,20) = 3.09$, $MS_e = 2626$, $p<.07$]. The pattern of times per database, however, is difficult to interpret with respect to the MVT; if participants assumed that longer processing delays made switching databases more costly, one would expect a pattern of times increasing from 2s to 4s and to 8s delay conditions.

As shown in Table 6, participants exhibited roughly the same GUTs, on average, in all delay conditions [$F(2,20) = 0.06$, $MS_e = 1646.1$, *n.s.*]. The GUN showed greater difference across conditions but these differences were not statistically significant [$F(2,20) = 0.66$, $MS_e = 7.80$, *n.s.*].

Overall, the results seem to indicate that participants did not show a consistent tendency to stay longer in databases when the database access delay was larger.

Table 5: *Measures of general foraging performance (Experiment 2).*

	Databases Opened		Items Opened		Items Selected		Time per Item (s)		Selected Item Rating (1–7)	
Delay	**Mean**	***S.D.***	**Mean**	***S.D.***	**Mean**	***S.D.***	**Mean**	***S.D.***	**Mean**	***S.D.***
2 s	4.67	*1.56*	81.17	*33.35*	22.67	*14.01*	16.61	*8.44*	5.32	*0.46*
4 s	5.67	*1.78*	66.08	*26.76*	18.58	*9.11*	19.25	*10.84*	5.43	*0.76*
8 s	5.17	*1.34*	52.75	*20.91*	14.83	*6.98*	25.00	*10.54*	5.25	*0.67*

S.D. Standard deviation

Table 6: *Measures of patch-leaving behaviour (Experiment 2).*

	Time in Database (s)		Items Opened per Database		Time in Database After Last Selection (s)		Items Opened in Database After Last Selection	
Delay	**Mean**	**S.D.**	**Mean**	**S.D.**	**Mean**	**S.D.**	**Mean**	**S.D.**
2 s	271.02	*77.17*	19.32	*10.61*	54.17	*67.88*	2.90	*3.97*
4 s	219.27	*70.60*	12.90	*7.85*	51.11	*41.90*	2.87	*3.85*
8 s	240.62	*68.75*	10.71	*5.29*	48.21	*27.22*	1.75	*1.23*

S.D. Standard deviation

4.2.3 Analysis question order

Table 7 shows data across analysis question order for the second experiment. There is little indication that participants' search behaviour changed over the course of the experimental session. Question order did not have a significant effect on the number of databases opened [$F(2,20) = 1.24$, $MS_e = 1.89$, *n.s.*], time spent in databases [$F(2,20) = 2.27$, $MS_e = 2626$, *n.s.*], time spent processing each item [$F(2,20) = 0.24$, $MS_e = 109.80$, *n.s.*], number of items opened per database [$F(2,20) = 1.77$, $MS_e = 23.11$, *n.s.*], or mean rating given to selected items [$F(2,20) = 1.07$, $MS_e = 0.32$, *n.s.*]. The only significant difference observed was on the number of items selected per database, with participants selecting more items in the first question than subsequent questions [$F(2,20) = 4.35$, $MS_e = 3.00$, $p<.05$].

Neither the GUT [$F(2,20) = 4.19$, $MS_e = 1646.41$, $p<.05$] nor GUN [$F(2,20) = 0.99$, $MS_e = 7.80$, *n.s.*] exhibited any significant difference from the first to third question.

4.2.4 Information gain rate at the time of leaving

Observed gain rates ATOL are shown in Table 8. No significant difference was found between delay conditions [$F(2,20) = 2.26$, $MS_e = 0.001$, *n.s.*].

Table 7: *Measures of foraging behaviour across analysis question order (Experiment 2).*

	Databases Opened		Items Opened per Database		Items Selected per Database		Time per Item (s)		Selected Item Rating (1–7)	
Question	**Mean**	**S.D.**	**Mean**	**S.D.**	**Mean**	**S.D.**	**Mean**	**S.D.**	**Mean**	**S.D.**
First	4.83	*1.64*	16.40	10.17	5.08	3.48	19.62	*9.10*	5.14	*0.56*
Second	5.67	*1.78*	12.90	7.85	3.65	2.52	19.25	*10.84*	5.43	*0.76*
Third	5.00	*1.20*	13.63	7.85	3.06	1.88	21.99	*11.16*	5.43	*0.54*

S.D. Standard deviation

Table 8: *Observed gain rate ATOL (Experiment 2).*

	Gain Rate At Time Of Leaving (Cumulative Rating per second)	
Delay	**Observed**	***S.D.***
2 s	0.101	*0.041*
4 s	0.094	*0.055*
8 s	0.069	*0.051*

S.D. Standard deviation

5 Conclusion

5.1 Summary

Unfortunately, the structure of the information environment in the two experiments presented only a minimally interesting problem for an information forager. Given the linear cumulative gain function for databases, the theoretical optimal stopping point for foraging within a database is at item selection number 40; i.e., a forager should forage a database until all items have been consumed. In practice, however, there is little difference in expected gain for foragers who depart a database at any point after the 6th or 7th selection because, at that point, the rate of gain is essentially flat.

In light of this situation, participants' behaviours in Experiments 1 and 2, while not technically optimal, are nevertheless quite rational. Participants tended to visit most databases, selecting on average close to 20 items before leaving. Certainly, participants' patch-leaving decisions were variable, with participants sometimes leaving a database after only a small number of selections or, other times, exhausting all items before leaving. Nevertheless, participants gathered roughly as much information overall all as could be expected.

Beyond this general finding that participants did not behave in any grossly maladaptive fashion, the results of the two experiments offer no indication that participants' foraging activities were sensitive to between- and within-patch costs. The extremely large degree of variability seen in all measures prevented any but a few measures from exhibiting statistically reliable differences. This variability could have resulted from several factors, including individual differences in reading speed and differences in foraging strategies. It may be possible to infer individual foraging strategies with larger numbers of participants. Nevertheless, on measures of general foraging behaviour and patch-leaving, participants seemed to behave similarly regardless of the between- or within-patch time costs attached.

The experiments presented a situation in which participants had access to no information about the quality of individual information items. This is fairly unrealistic condition as information systems are generally designed to help users find whatever information they seek. Thus, items and links are labelled and often summaries or representative sections can be viewed so that users can estimate the value of information before making a navigation decision. This sort of navigation support was not provided specifically to allow examination of participants' foraging judgments when based solely on what they were able to learn about the statistical distribution of information in the environment. Although it appears that participants were able to comprehend the basic nature of the information environment and forage effectively overall, they did not show any indication of learning the specific distributions of relevant items across databases within an experiment. Despite foraging in three analysis questions in which the ratios of relevant to irrelevant information items across databases, participants did not show any sign of exploiting knowledge of this fact. They did not bias their foraging to richer databases in the second or third question nor performed better overall on the third versus first question. In light of this finding, intelligence analysts would likely benefit from some form of decision support that could highlight regularities in the distribution of relevant information across available information sources.

5.2 Limitations

The most serious limitation of the current research was the specification of the databases in the experiment to exhibit linear cumulative gain functions and, hence, a constant rate of gain. This violated an assumption of Charnov's MVT, substantially weakening the potential for the experiments to provide evidence concerning the question of whether peoples' foraging behaviour conforms to predictions of IFT. In designing the experiments, it was assumed that the databases would exhibit diminishing returns because the probability of selecting a relevant item would decrease with each selection. This was true but neglected the true cumulative value expected to accumulate with each selection (as demonstrated in the Predictions section). In particular, the expected value of items remained constant as participants selected items. In natural environments, however, the expected value of items will tend to decrease as a forager exploits a patch (because a forager will generally selected the best looking items first). In addition, the time delay on opening individual items was held constant within each analysis question. In natural environments, the handling time needed for items may increase as a forager exploits a patch (because items become more physically dispersed over time, for example), further diminishing the rate of gain in that patch. Subsequent research will correct for these flaws in the experimental design.

A second major limitation was the failure to secure a sufficient number of trained intelligence analysts as participants. This meant it was impossible to contrast the foraging performance of experts and novices in the simulated analysis task.

5.3 Future research directions

The current research was a tentative first step and a great deal more remains to be done on the topic of IFT. Future work can be organized around a few broad questions, as discussed below.

5.3.1 Do people conform to predictions of IFT?

The main issue addressed by the current research was whether, or to what extent, peoples' information search behaviour conforms to predictions of IFT. The results offered no evidence that participants were sensitive to between- and within-patch costs. As discussed, however, the two experiments presented a very constrained information environment in which patch cumulative gain functions did not support a rigorous test of the predictions of the MVT. Before concluding that people do not employ optimal foraging strategies, additional studies must examine a wider range of information foraging tasks and environments. In particular, these experiments should be re-run with databases in which the cumulative gain functions do exhibit a pattern of diminishing returns.

Subsequent experiments should ensure that participants forage in environments in which each patch exhibits a pattern of diminishing returns with each item selection, as assumed by Charnov's MVT [18]. Having participants forage in environments that meet this assumption poses a real test of peoples' sensitivity to between- and within-patch costs. A pattern of diminishing returns could be achieved by altering the item processing cost in such a way that the amount of time required to open and process each item increases as a function of time within the database. Conforming to optimal foraging predictions in this case would be a more significant indicator as to whether people are able to make adaptive foraging decisions in an information space.

Although not considered in the current experiments, IFT also allows predictions pertaining to item selection. This is referred to as the "diet problem," which, in the context of information foraging, deals with a forager's criterion for accepting or rejecting a piece of information. In particular, IFT makes several general predictions about how strict a forager will make his or her criterion for acceptance, depending on the environmental conditions. One such prediction has to do with the impact of opportunity costs on items selection. The processing time required to exploit an item affects that item's relative profitability. Generally, longer processing times result in greater opportunity costs because the time spent processing an item is necessarily time that cannot be used to search for other items [10, pp. 41–42]. This means that, all else being equal, difficult-to-process items are less attractive than items that require less effort and time.

Another prediction is that the prevalence, or frequency of encounter, of a given type of item does not determine its attractiveness to a forager but, instead, the relative prevalence of more valuable types of items. This is the principle of *Independence of Inclusion from Encounter Rate,* which states that the decision to pursue a class of items is independent of its own prevalence in the environment [11]. Based on this prediction, we would expect an information forager to alter his or her criterion for accepting information items [10, pp. 41–42]. In a rich environment (i.e., an environment containing many, highly rated relevant items), a forager should be more selective, which would be reflected in choosing items that have a higher relevance rating on average. In a poor environment, with few and/or lower-rated items, a forager should be less selective and accept items with lower relevance ratings on average.

In the context of an INFOCAT experiment, this principle can be tested by varying the ratio of very highly rated items to lower rated items. It would be predicted that the number of low-rated items in databases would not itself affect a forager's criterion for accepting an item. Thus, if databases contained more items of low quality, a participant would not lower his or her criterion to accept such items. If, however, the frequency of high-rated items was increased relative to low-rated items, foragers would be expected to change their criterion; in this case raising it to reject lower-rated items that would have been acceptable before.

5.3.2 Do people make effective use of information "scent" cues?

Information scent refers to the process by which information foragers make use of contextual cues to enhance their information foraging performance [12]. It is analogous to the sensory cues animals use to enhance their predictions of the location and quality of foods. Thus, cues associated with navigation options that provide a forager with some indication of the nature and value of the content available by these options [10, p.68][15]. Unlike sensory cues, however, information scent generally refers to semantic context associated with navigation options for getting information.

Information scent has become an important concept for understanding peoples' information search decisions and as a way to enhance the usability of information search tools. Almost universally, web pages contain labeled navigation links from one web page to another and Web page designs continue to evolve to more effect scent cues, such as snippets of text and graphics with such links [10, p. 69].

One question is how information scent is assessed by information foragers. It is assumed that an information forager evaluates items by a semantic comparison process in which the identifier for the patch (e.g., title, link descriptor, etc.) is matched to the forager's internal representation of the information goal [15]. The overlap between the meaning of the search goal and patch identifier is the basis for a quantitative measure of similarity, which is viewed as a proxy measure for the expected value of accessing the patch. In practice, there are multiple ways to conceptualize and measure semantic similarity, including objective semantic analysis techniques, such as Latent Semantic Similarity (LSA) [14][25].

The current research purposefully avoided providing scent cues in order to allow us to study how foragers' behaviour was governed by their appreciation of the statistical structure of the information environment. Nevertheless, future research is needed to examine the extent to which information foragers make use of information scent to guide their search and what aspects of scent cues are most effective in helping people forage.

5.3.3 Do people adapt their foraging strategies to the environment?

The current experiments yielded no support for the hypothesis that participants could learn to make use of regularities in the statistical distribution of relevant information items across databases. The experimental design, however, employed just three analysis questions and it may have been difficult for participants to notice differences in the relative richness of the databases. Thus, it is too soon to rule out the possibility that people could adapt to specific environmental conditions with more experience.

If information foragers are able to derive an understanding of how information is distributed across patches in the information environment, they may be able adopt simple heuristic rules for patch-leaving that will allow optimal foraging performance (see [22]). Green [17][26] has demonstrated optimal patch-leaving rules for environments with certain general distribution types. In brief, when items are uniformly distributed across patches (i.e., patches have the same number of items) the optimal strategy is the Fixed-Number rule, in which a forager sets a criterion number of items to exploit in each patch then leaves a patch when that number has been obtained. When items are distributed across patches according to a Poisson distribution, the Fixed-Time rule, in which a forager sets a fixed amount of time to spend in each patch, regardless of his/her success or failure, is optimal. When variability in the richness of patches is greater, and the environment is best described by a negative-binomial distribution, the optimal strategy is the incremental rule [27]. This heuristic works by setting an amount of time to be spent searching a patch. Then, as that time is expended, additional time is added when an item is successfully exploited.

The possible use of these sorts of simple heuristics could be explored in INFOCAT experiments in which the distribution of relevant items across databases is manipulated. With a sufficiently large number of databases, the uniform, Poisson, and negative-binomial distributions could be approximated, producing environments in which different patch-leaving rules are predicted to be optimal. If participants are sensitive to the underlying distribution of relevant items across databases, and are able to adaptively alter their foraging behaviour, their foraging performance (as indicated by various measures) should differ across conditions. This approach has been applied by Louâpre, van Alphen, and Pierre [28], although their results suggest people are not able to adaptively switch patch-leaving strategies in response to different environmental distributions.

References

[1] Rudner, M. (2002). The future of Canada's defence intelligence. *International Journal of Intelligence and CounterIntelligence, 15*, 540–564.

[2] Russell, D. M., Stefik, M. J., Pirolli, P., & Card, S. K. (1993, May). The cost structure of sensemaking. In *Proceedings of the INTERACT'93 and CHI'93 conference on Human factors in computing systems* (pp. 269–276). ACM.

[3] Pirolli, P., & Card, S. (2005, May). The sensemaking process and leverage points for analyst technology as identified through cognitive task analysis. In *Proceedings of International Conference on Intelligence Analysis* (Vol. 5, pp. 2–4).

[4] Badalamente, R. V., & Greitzer, F. L. (2005, May). Top ten needs for intelligence analysis tool development. In *proceedings of the 2005 international conference on intelligence analysis*.

[5] Puvathingal, B. J., & Hantula, D. A. (2012). Revisiting the psychology of intelligence analysis: From rational actors to adaptive thinkers. *American Psychologist, 67*, 199–210.

[6] Pfautz, J., Fichtl, T., Guarino, S., Carlson, E., Powell, G., & Roth, E. (2006). *Cognitive Complexities Impacting Army Intelligence Analysis*. Proceedings of the Human Factors and Ergonomics Society 50th Annual Meeting (pp. 452–456). Sage.

[7] Hutchins, S. G., Pirolli, P., & Card, S. (2003, October). *Use of critical analysis method to conduct a cognitive task analysis of intelligence analysts*. Proceedings of the Human Factors and Ergonomics Society Annual Meeting (Vol. 47, No. 3, pp. 478–482). SAGE Publications.

[8] Cook, M. B., & Smallman, H. S. (2008). Human factors of the confirmation bias in intelligence analysis: Decision support from graphical evidence. *Human Factors, 50*, 745–754.

[9] Bryant, D. J. (2014). *Information Foraging Theory: A framework for intelligence analysis*. DRDC Scientific Report (DRDC-RDDC-2014-R115). Defence Research and Development Canada – Toronto Research Centre.

[10] Pirolli, P. (2009). Information foraging theory: Adaptive interaction with information. New York, NY: Oxford University Press.

[11] Pirolli, P., & Card, S. (1999). Information foraging. *Psychological review*, *106*(4), 643.

[12] Chi, E. H., Pirolli, P., Chen, K., & Pitkow, J. (2001). *Using information scent to model user information needs and actions on the web*. SIGCHI'01, Seattle, WA.

[13] Mantovani, G. (2001). The Psychological Construction of the Internet: From Information Foraging to Social Gathering to Cultural Mediation. *CyberPsychology & Behavior, 4*, 47–56.

[14] Budiu, R., Royer, C., & Pirolli, P. (2007, May). Modeling information scent: A comparison of LSA, PMI and GLSA similarity measures on common tests and corpora. In *Large Scale Semantic Access to Content (Text, Image, Video, and Sound)* (pp. 314–332). Le Centre de Hautes Estudes Internationales d`ìnformatique documentaire.

[15] Pirolli, P. (2005). Rational analysis of information foraging on the web. *Cognitive Science, 29*, 343–373.

[16] Payne, S. J., Duggan, G. B., & Neth, H. (2007). Discretionary Task Interleaving: Heuristics for Time Allocation in Cognitive Foraging. *Journal of Experimental Psychology: General, 136*, 370–388.

[17] Green, R. F. (1988). *Optimal foraging for patchily distributed prey: random search*. Technical Report 88-2. University of Minnesota. Department of Mathematics and Statistics.

[18] Charnov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theoretical population biology*, *9*, 129–136.

[19] Wilke, A., Hutchinson, J. M. C., Todd, P. M., & Czienskowski, U. (2009). Fishing for the Right Words: Decision Rules for Human Foraging Behavior in Internal Search Tasks. *Cognitive Science, 33*, 497–529.

[20] Mobus, G. E. (1999). *Foraging Search: Prototypical Intelligence*. Paper presented at The Third International Conference on Computing Anticipatory Systems, Liege, Belgium.

[21] Stephens, D. W., & Krebs, J. R. (1986). *Foraging theory*. Princeton, NJ: Princeton University Press.

[22] Bryant, D. J. (2017). *Simple heuristics for guiding information search.* DRDC Scientific Report (DRDC-RDDC-2017-R119). Defence Research and Development Canada – Toronto Research Centre.

[23] Puvathingal, B. J. (2013). *Homo informaticus intelligens: Building a theory of intelligence analysts as information foragers*. Unpublished Doctoral Dissertation. Temple University, Department of Psychology. UMI 3564770.

[24] Bryant, D. J., & Li, A. (2016). INformation FOraging Cognitive Analysis Tool (INFOCAT): An experimental platform for studying information foraging of intelligence analysts. DRDC Scientific Report (DRDC-RDDC-2016-R022). Defence Research and Development Canada – Toronto Research Centre.

[25] Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review, 104*, 211–240.

[26] Green, R. F. (1979). *Bayesian birds: a simple example of Oaten's stochastic model of optimal foraging*. Technical Report No. 50. Department of Statistics, University of California, Riverside.

[27] Hansjörg, N., Schooler, L., Rieskamp, J., Quesada, J., Xiang, J., Wang, R., Wang, L., Zhou, H., Qin, Y., Zhong, N., & Zeng, Y. (2009). Analysis of Human Search Strategies. *Analysis*, *4*, 2.

[28] Louâpre, P., van Alphen, J. J., & Pierre, J. S. (2010). Humans and insects decide in similar ways. *PLoS One*, *5*(12), e14251.

List of symbols/abbreviations/acronyms/initialisms

ATOL	at the time of leaving
DRDC	Defence Research and Development Canada
GUN	Giving-Up-Number
GUT	Giving-Up-Time
IFT	Information Foraging Theory
INFOCAT	INformation FOraging Cognitive Analysis Tool
JICAC	Joint Intelligence Collection and Analysis Capability
LSA	Latent Semantic Similarity
MSe	Mean Square Error
MVT	Marginal Value Theorem
n.s.	not significant
p	probability
s	second
S.D.	Standard Deviation
USS	United States Ship

DOCUMENT CONTROL DATA

(Security markings for the title, abstract and indexing annotation must be entered when the document is Classified or Designated)

1. ORIGINATOR (The name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g., Centre sponsoring a contractor's report, or tasking agency, are entered in Section 8.) DRDC – Toronto Research Centre Defence Research and Development Canada 1133 Sheppard Avenue West P.O. Box 2000 Toronto, Ontario M3M 3B9 Canada	2a. SECURITY MARKING (Overall security marking of the document including special supplemental markings if applicable.) CAN UNCLASSIFIED 2b. CONTROLLED GOODS (NON-CONTROLLED GOODS) DMC A REVIEW: GCEC DECEMBER 2013

3\. TITLE (The complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S, C or U) in parentheses after the title.)

Experiments on Information Foraging

4\. AUTHORS (last name, followed by initials – ranks, titles, etc., not to be used)

Bryant, D.J.

5. DATE OF PUBLICATION (Month and year of publication of document.)	6a. NO. OF PAGES (Total containing information, including Annexes, Appendices, etc.)	6b. NO. OF REFS (Total cited in document.)
November 2017	40	28

7\. DESCRIPTIVE NOTES (The category of the document, e.g., technical report, technical note or memorandum. If appropriate, enter the type of report, e.g., interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.)

Scientific Report

8\. SPONSORING ACTIVITY (The name of the department project office or laboratory sponsoring the research and development – include address.)

DRDC – Toronto Research Centre
Defence Research and Development Canada
1133 Sheppard Avenue West
P.O. Box 2000
Toronto, Ontario M3M 3B9
Canada

9a. PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.)	9b. CONTRACT NO. (If appropriate, the applicable number under which the document was written.)
JFD 4.1	

10a. ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document.)	10b. OTHER DOCUMENT NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)
DRDC-RDDC-2017-R159	

11\. DOCUMENT AVAILABILITY (Any limitations on further dissemination of the document, other than those imposed by security classification.)

Unlimited

12\. DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11) is possible, a wider announcement audience may be selected.))

Unlimited

13. ABSTRACT (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

Two experiments tested predictions of Information Foraging Theory (IFT) pertaining to "patch-leaving," the decision to abandon an information source to search elsewhere. IFT predicts that increasing the between-patch cost associated with moving from one source of information to another should lead foragers to increase the time spent in each patch. Similarly, increasing the within-patch cost associated with processing individual information items should likewise increase the time spent in each patch. Participants searched for information relevant to solving a series of simulated analysis questions using the INformation FOraging Cognitive Analysis Tool (INFOCAT) platform. Information items were separated into a number of discrete databases and participants were allowed to freely search information and select items they judged to be relevant. In Experiment 1, the time delay associated with opening a database (between-patch time) was varied across analysis questions, while in Experiment 2, the time delay added to opening an information item (within-patch time) was varied across analysis questions. The results of the two experiments indicated no evidence that participants' patch-leaving decisions were affected by either between- or within-patch costs. There was also no indication that participants' search behaviour changed over the course of an experimental session. These results suggest that people do not necessarily apply optimal foraging strategies to information search. Although further research is needed to conclusively determine whether peoples' information foraging conforms to predictions of IFT, an opportunity exists to enhance the efficiency and effectiveness of the intelligence analysis process by translating IFT concepts into training and decision support for intelligence analysts.

Deux expériences ont vérifié les prédictions de la théorie du butinage des renseignements (TBR) concernant la décision d'abandonner les recherches dans une source d'information pour poursuivre les recherches ailleurs. La TBR prédit que l'augmentation du coût associé au déplacement d'une source d'information à une autre devrait inciter les butineurs à passer plus de temps dans chaque source d'information. Dans un même ordre d'idées, l'augmentation du coût associé au traitement d'éléments d'information individuels devrait aussi inciter les butineurs à passer plus de temps dans chaque source d'information. Les participants ont fait des recherches d'information afin de répondre à une série de questions d'analyse simulées à l'aide de l'outil d'analyse cognitive du butinage de renseignements (INFOCAT). Les éléments d'information ont été divisés en un certain nombre de bases de données distinctes. Les participants pouvaient mener leurs recherches librement et choisir les éléments qu'ils jugeaient pertinents. Dans l'expérience 1, le délai d'ouverture d'une base de données (pendant la transition d'une source à une autre) variait selon les questions d'analyse, alors que dans l'expérience 2, le délai d'ouverture d'un élément d'information (pendant la recherche dans une source) variait selon les questions d'analyse. Les résultats des deux expériences n'ont donné aucune preuve que le coût de la transition ou de la recherche avait une incidence sur la décision de passer à une nouvelle source d'information. Rien n'indique non plus que les participants ont modifié leur comportement de recherche au cours des deux expériences. Ces résultats portent à croire que les gens n'appliquent pas nécessairement des stratégies de butinage des renseignements optimales. Bien que d'autres recherches soient nécessaires pour établir avec certitude que le butinage de renseignement est conforme aux prédictions de la TBR, il est possible d'améliorer l'efficacité

du processus d'analyse du renseignement en convertissant les concepts de la TBR en formation et en aide aux décisions pour les analystes du renseignement.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g., Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Information foraging; heuristic; optimization