Hybrid value foraging: How the value of targets shapes human foraging behavior

Jeremy M. Wolfe^{1,2} · Matthew S. Cain^{2,3} · Abia Alaoui-Socé²

Published online: 21 December 2017
© The Psychonomic Society, Inc. 2017

Abstract

In hybrid foraging, observers search visual displays for multiple instances of multiple target types. In previous hybrid foraging experiments, although there were multiple types of target, all instances of all targets had the same value. Under such conditions, behavior was well described by the marginal value theorem (MVT). Foragers left the current “patch” for the next patch when the instantaneous rate of collection dropped below their average rate of collection. An observer’s specific target selections were shaped by previous target selections. Observers were biased toward picking another instance of the same target. In the present work, observers forage for instances of four target types whose value and prevalence can vary. If value is kept constant and prevalence manipulated, participants consistently show a preference for the most common targets. Patch-leaving behavior follows MVT. When value is manipulated, observers favor more valuable targets, though individual foraging strategies become more diverse, with some observers favoring the most valuable target types very strongly, sometimes moving to the next patch without collecting any of the less valuable targets.

Keywords visual search · Attention: Selective · attention and memory

Imagine searching a collection of coins for all U.S. quarters (\$0.25), dimes (\$0.10), nickels (\$0.05), and pennies (\$0.01). This is an example of a hybrid foraging task. “Hybrid” search tasks are searches for any of several possible targets. Hybrid foraging tasks are searches for multiple instances of several possible targets. How are such tasks influenced by the relative value of the targets and by the relative prevalence of those targets (e.g., should you search for rare quarters or for more common nickels)? This work extends the search literature to new tasks that occur regularly in the world beyond the laboratory.

Our days are filled with searches: Sifting through Internet results to find the webpages we want (Pirulli, 2007), looking through the trail mix to find our favorite nuts, combing through a child’s hair to find elusive nits and lice. Notice that

each of these examples are complex searches that cannot be described as simple two-alternative, forced-choice, “present” versus “absent” search tasks of the sort most typically conducted in lab settings (for recent reviews, see Chan & Hayward, 2012; Wolfe, 2014a, b). Rather, these searches are “hybrid foraging” tasks. Hybrid foraging combines the characteristics of “hybrid search” and “foraging” tasks. A hybrid search task is a visual search for an instance of any of several possible targets held in memory (Schneider & Shiffrin, 1977; Wolfe, 2012). A foraging task is a search for multiple, generally unknown number of instances of one type of target (Bond, 1981; Cain, Vul, Clark, & Mitroff, 2012; Stephens & Krebs, 1986; Wolfe, 2013). Thus, hybrid foraging tasks are searches for multiple visual instances of several target types, held in memory. Such tasks involve search through both our memories and the visual displays presented to us (Schneider & Shiffrin, 1977).

✉ Jeremy M. Wolfe
jwolfe@bwh.harvard.edu

¹ Harvard Medical School, Boston, MA, USA

² Visual Attention Lab, Brigham & Women’s Hospital, Boston, MA, USA

³ U.S. Army Natick Soldier Research, Development, and Engineering Center, Natick, MA, USA

The Hybrid Foraging Paradigm

Kristjansson, Johannesson, and Thornton (2014) described a version of a hybrid foraging task. In their task, observers collected all examples of two types of targets in a display. Thus, observers might be asked to pick red and green items in a

display of red, green, blue, and yellow items. Observers were required to pick all target items in a display. Kristjansson and colleagues were most interested in the sequence of selections. Did observers pick at random between the two target types, or did they tend to pick targets in runs—a set of red items, then a set of green, and so forth. The answer depended on the task. In the easy color search, described here, observers tended to switch back and forth frequently. In a more difficult conjunction search, they tended to pick one type of item for a while before switching to another. Indeed, they generally picked all of one item before switching.

Wolfe, Aizenman, Boettcher, & Cain (2016) used a somewhat different hybrid foraging paradigm. Participants searched moving displays of many objects for instances of any of several target objects held in memory. In these experiments, the memory set size varied from eight to 64 in different blocks. Objects moved about the screen in order to thwart “reading” strategies in which observers simply started at the upper left and searched to the lower right of the display. In this article, following the jargon of the animal foraging literature, each screenful of objects will be called a “patch.” A session in which observers move through a series of patches is a “block.” Unlike Kristjansson et al. (2014), observers in the Wolfe et al. (2016) study moved to the next patch whenever they wanted, allowing us to study patch-leaving times, a variable of interest in the foraging literature. Like Kristjansson et al. (2014), Wolfe et al. (2016) found that the identity of the most recently selected target in a patch is biased toward the previously selected target type. That is, if an observer was searching for red apples, yellow trucks, and white washing machines, picking a truck made it more likely that the next selection would also be a truck. Thus, targets were collected in “runs” that were longer than those predicted by random selection among available targets. Such runs would be predicted if one assumes that finding a target primes the features of that type of target (Kristjansson, 2006; Maljkovic & Nakayama, 1994; Olivers & Hickey, 2010). Finding a yellow truck would bias subsequent search toward other items that were yellow (and that had other truck-like basic features). Run behavior was less dramatic than in some Kristjansson et al. (2014) conditions where observers seem to make a strategic decision to collect all the instances of one target type before moving to the other type.

In foraging tasks, where the number of targets is unknown, it is important to know when the observer chooses to leave the current patch to “travel” to the next (Bond, 1981; Pyke, Pulliam, & Charnov, 1977). This is not an issue if observers are required to collect every target. However, it is important in tasks where the forager is free to move to a new patch at will. In an earlier foraging study with a single target type, Wolfe (2013) found that participants behaved in a manner that generally followed Charnov’s (1976) marginal value theorem (MVT), which describes patch-leaving behavior in animal foragers, at least in simple foraging situations. MVT states that

foragers will leave a patch for a new one when the “instantaneous rate of return” from the current patch drops below the average rate of return over all patches. The instantaneous rate of return when all items are of equal value is the inverse of the average response time (RT) for that moment. If it takes an average of 500 ms to collect the fifth target in a patch, the instantaneous rate of return would be 2.0 targets a second (assuming that all of your clicks fall on actual targets). As noted, average patch leaving behavior in hybrid foraging was also well-described by MVT (Wolfe et al., 2016).

In the present work, we are interested in the effect of target value on hybrid foraging behavior. In the Wolfe et al. (2016) hybrid foraging experiment, all targets were equally valuable. Participants accumulated the same number of points for each target collected. However, in real-world foraging tasks, targets can have different values. Thus, you might look for both peanuts and cashews in the trail mix, but of the two, you might prefer cashews. Cashews would be the more valuable search target. How does our foraging behavior adapt to searches for differently valued targets? Previous studies have repeatedly demonstrated that reward modulates how attention is deployed toward selected, behaviorally relevant items (for an extensive review, see Failing & Theeuwes, 2017). The prospect of greater reward biases attention toward specific stimuli (e.g., Della Libera & Chelazzi, 2006; Navalpakkam, Koch, Rangel, & Perona, 2010; Serences, 2008;). Valuable items serve as more effective cues for attention than less valuable items (Munneke, Hoppenbrouwers, & Theeuwes, 2015). Added value can behave as though it is adding to the relative salience of targets (Hickey, Chelazzi, & Theeuwes, 2010). Navalpakkam et al. (2010) looked at the comparative influence of value and salience on search for multiple targets in a complex perceptual environment. They found that the decisions made by searchers accounted for both value and salience in a manner consistent with the ideal (Bayesian) combination of these cues. In the present work, value is manipulated only by assigning different numbers of points to different target types. Collecting more “valuable” targets allowed observers to finish the task more quickly and, presumably, were rewarding in the same way that video game points are valuable. The value manipulation produces significant effects—effects that one might imagine would be stronger if the point values translated into a more concrete value like money.

In the real world, valuable items are likely to be rarer than less valuable items. In a search task, this is interesting because, while value makes items more attractive to attention, low prevalence makes them more likely to be overlooked (Bond & Kamil, 2002; Wolfe, Horowitz, & Kenner, 2005). Thus, it is important to understand how prevalence and value interact in a hybrid foraging situation. What strategy do you adopt if you are searching for common, low-value peanuts in a bowl of trail mix that may contain rarer but higher value cashews? In the

current experiment, we ask how the prevalence and value of targets interact to influence foraging strategies and attentional deployment.

Method

Participants

Twelve naïve observers (nine females), ages 18 to 47 years ($M = 24.42$ years, $SD = 7.86$ years) participated in the experiment. All participants had normal or corrected-to-normal vision and passed the Ishihara Color Test (Ishihara, 1980). They provided oral informed consent and received a total of \$15 for this experiment, which typically took an hour and a half to complete. The procedures employed here all were approved by the Partners Healthcare Corporation Institutional Review Board.

Stimuli and apparatus

In this hybrid foraging experiment, observers searched for multiple instances of multiple targets. As in Wolfe et al. (2016), observers collected targets from a succession of visual displays (patches). The patches contained 60, 75, 90, or 105 items (visual set size), programmed to move continuously in random directions at a rate of 1.25°/s in order to discourage a reading strategy. Items followed randomly defined trajectories. They were repulsed by the edges and center of the display, as well as by other items. There were different levels of repulsion exerted by the edges, the center, and by other objects. These forces interacted in a manner that allowed some overlap of objects when the objects' "desire" not to overlap was overcome by the stronger repulsion from elsewhere (e.g., the edge of the display). Clicking on a target item caused that item to vanish and gave points to the observer. Clicking on a distractor also caused that item to vanish but deducted a point from the observer. The goal was to accumulate points to reach a prespecified total number of points as quickly as possible. Measures of interest included the observers' rates of collection, their choice of targets, and the time at which they chose to leave each patch in order to start searching in a new one.

The experiment was written in MATLAB 8.3 (MathWorks, Natick, MA) using Version 3.0 of the Psychophysics Toolbox (Brainard, 1997; Pelli, 1997). The stimuli were photographs of objects drawn from the 1,314 item set used in Brady, Konkle, Alvarez, and Oliva's (2008) picture memory experiments. They were presented on a 24in. LCD monitor (Mitsubishi Diamond Pro 91TXM) set at a 1920 × 1200-pixel resolution, with a 60-Hz refresh rate. Observers were placed at a viewing distance of approximately 57 cm, such that 1 cm subtended approximately 1 degree of visual angle. Consequently, the display subtended 52 × 32 degrees. Observers were free to move their eyes.

Procedure

In each of three conditions, observers held four targets in memory, learning a new set of four targets for each condition. Targets to be memorized were picked at random from all of the available objects. The design used for the visual displays or "patches" in this experiment closely resembled the one used in a previous hybrid foraging experiment (Wolfe et al., 2016; see Fig. 1 for an example). Each patch started with 60, 75, 90, or 105 total items. Of these, 20%–30% of all items were targets before the observer began collecting targets and removing them from the display. The remaining 70%–80% of items were distractors. The distractors were chosen so that the number of instances of any specific distractor type was neither markedly larger nor smaller than the average number of instances of the different target types. Thus, since about one fourth of items were targets, the three fourths of items that were distractors were divided among roughly nine types of distractor in each patch. Observers completed three different conditions that differed in the point value and prevalence of the four targets held in memory, as described below.

To begin each condition, participants memorized the set of four targets. Each was shown for 3 seconds, along with its associated point value so observers knew the values from the start of a block and did not need to learn them while foraging. Observers were then tested for target recognition. Eight images were presented one at a time. Four were the targets and four were foils. Participants made a forced choice target/nontarget response for each. If participants made any errors, they repeated the memorization and recognition test sections.

Participants collected items by clicking on them, receiving points for correct (target) collections and losing a point for incorrect (distractor) collections. A score report was displayed at the center of the screen in black, turning red whenever points were lost. To move on to a new patch, participants could click the "next" button at any time, also located at the center of the screen. The time between patches is known as the "travel time" in the animal literature where, after all, the animal is travelling. Here, there was a programmed 2-second minimum travel time. The actual (because of the time required to calculate the next display) measured travel time between the last collection in a patch and the appearance of the next patch averaged approximately 5 seconds. Travel time was not manipulated as an independent variable in this experiment. If it had been, longer travel time would be expected to induce observers to stay longer in the current patch (though human have complex responses to manipulations of travel time in foraging tasks; see Wolfe, 2013). For each of the three conditions, participants completed a practice and experimental block. After each patch in the practice blocks, missed targets were outlined with boxes on the screen as error feedback for participants. No miss feedback was provided during the experimental blocks.

Fig. 1 Cartoon of a hybrid foraging display. Four targets are shown at the top. These would not be on screen during the trial. Observers would memorize the targets and their associated values. Then they would click on examples in the display. Actual displays would contain more and

different items, and each item would move randomly. Observers would click on a “next” button at will to move to a new screen. (Color figure online)

As noted above, participants completed three conditions, differentiated by the value and prevalence of the four, memorized search targets. Participants learned four new targets for each condition. The order of conditions was counterbalanced across participants. On average, participants took about 16 minutes to complete each condition. We will refer to the value across targets as being “even” or “uneven” and to the prevalence of different targets as “equal” or “unequal.” We could use the same terms (e.g., equal/unequal) in both cases, but we hope the different terms help to differentiate the conditions.

Even value, unequal prevalence In this condition, the value of each of the four targets was set at 4 points. However, the prevalence of each target was varied: Of the four targets, 53% were of the first target type, 27% were of the second target type, 13% of the third type, and 7% of the fourth type. To finish the practice block, participants had to collect 200 points, and to finish the experimental block, participants had to collect 4,000 points. (Note: point values were chosen in order to create blocks that required about 15 minutes to complete for an average observer.)

Uneven value, equal prevalence In this condition, the prevalence of all four targets was set at 25% while the value of each target was varied. One of the four target types (arbitrarily chosen) was set to be worth 2 points, a second was worth 4 points, a third was worth 8 points, and a fourth was worth 16 points. To finish the practice block, participants had to reach 400 points, and to finish the experimental block, participants had to reach 8,000 points. Again, the specific goal of 8,000 points was set to yield about 15 minutes of data. We assume that the “value” of a target is assigned relative to other targets within a block and that the way to think about the value of two targets would be to note, for example, that Target 1 is worth half of what Target 2 is worth (not that it is worth 4 points less. This assumption is not tested in these experiments.)

Uneven value, unequal prevalence In this condition, both the value and prevalence of the targets were varied. Prevalence was inversely related to value. A first target was worth 2 points and appeared with 53% frequency; a second target was worth 4 points and appeared with 27% frequency; a third target was worth 8 points, but appeared with 13% frequency;

and a fourth target was worth 16 points and appeared with 7% frequency. Note that, in this condition, the summed value of all instances of one type of target is the same as the summed value of any other type of target. That is, if an observer collects all of the 4-point objects, they would score as many points as if they collected all of the scarcer but more valuable 16-point objects. To finish the practice block in this condition, participants had to reach 250 points, and to finish the experimental block, participants had to reach 5,000 points.

Results

What do observers pick?

The hybrid foraging paradigm produces data that can be examined in many ways. One useful way to summarize the results is shown in Fig. 2.

In Fig. 2, the solid lines show the percentage of selections for each of the four target types as a function of the order of those selections (clicks) within a patch. The dotted lines show the relative percentages of each target type. This changes over the course of selections in a patch since selection of one type of target necessarily reduces the relative percentage of those targets in the display. For both proportion picked and proportion on-screen, the data are the average of individual observer averages. That is, the average percentage at each click is calculated for each observer. Those averages are, themselves, averaged, and those results are plotted with error bars representing ± 1 SEM. Data are plotted only for click positions with data from all 12 observers. This assures that every observer contributes to every data point. It does not mean they contribute equally. An observer who only reaches, say, 20 clicks in three patches in one condition makes the same contribution as an observer who reaches 20 clicks in 30 patches. As can be seen from the different numbers of clicks plotted in the three panels of Fig. 2, observers tend to stay in a patch for different lengths of time depending on the condition. It appears that observers leave the patches more readily when the targets are of different values than when they are of even value but unequal prevalence. Figure 3 gives a feeling for the variability in behavior. It shows the proportion of patches that receive N clicks. Of course, all patches receive at least one click and the proportion falls off thereafter. In the even-value, unequal-prevalence condition, all observers do something very similar. They collect an average of 18.0 items ($SD = 5.0$). In the uneven-value, equal-prevalence condition, behavior varies more across and within observers as can be seen by the shallow slope of the function and the larger error bars. The also pick fewer items (Avg. = 15.3, $SD = 6.0$). Finally, in the uneven-value, unequal-prevalence condition, behavior is even more variable (Avg. = 13.1, $SD = 7.5$) with one observer staying only for 4.3 clicks in an average patch and the most

Fig. 2 What do observers pick? For each of the three conditions, the figure shows the proportion of choices of each of the four targets as a function of time/clicks in a patch (solid lines). Dotted lines show the actual percentage of targets remaining of each of the four types for each click. Data are plotted for clicks where all 12 observers are represented. Error bars are ± 1 SEM. (Color figure online)

Fig. 3 Proportion of patches receiving at least N clicks. Data points are averaged over 12 observers. Error bars show ± 1 SEM for that average. (Color figure online)

persistent observer staying for 20.3 clicks. We will return to this topic later.

Returning to Fig. 2: If observers were simply picking at random among the four target types, then the percentages picked would mirror the percentages in the display. Thus, if half the targets were Type 1, half the selections would be Type 1 and the solid and dotted lines of Fig. 2 would lie, more or less, on top of each other. Clearly this is not the case. Prevalence and value have effects on the selections. These can be seen by considering each of the conditions in turn.

Even value, unequal prevalence (Fig. 2a) In this condition, with no variation in value, we can see the pure effects of prevalence. As would be predicted, observers pick the common items at a higher rate than the less common items. Interestingly, observers tend to overpick the most prevalent target and underpick the two least prevalent targets. Targets are numbered and colored according to prevalence: 1 (red) being the most prevalent and 4 (green) being the least prevalent. To test the hypothesis that observers are picking in proportion to the prevalence of the items in the display, we average each observer's data over the first five selections. We use the first five because the earlier selections in a patch occur before those selections have, themselves, markedly altered overall prevalence. The specific choice of five is otherwise arbitrary. The picking rate and the percentage in the display are clearly not independent variables. If the observer preferentially picks Target 1, its percentage in the display declines. For this reason, we are looking only at the first five clicks in a patch and *pretending* that the overall percentages of different targets do not change over those five clicks. This is not strictly true, but the deviations are small enough that it seems legitimate to treat target type and picked versus actual percent as independent variables. A two-way ANOVA with those variables reveals an unsurprisingly large effect of target type, $F(3, 33) = 416, p < .0001$, generalized eta-squared (η_G^2) = .93. There is no effect of picked versus actual, $F(1, 11) = 0.67, p = .67, \eta_G^2 = 0.0$, because the effects go in different directions for different target types. The most prevalent items are overpicked relative to their prevalence in the display. The least prevalent are underpicked. Thus, the interaction is significant, $F(3, 33) = 6.3, p = .0016, \eta_G^2 = .27$. Post-hoc t-tests show that Target 1 is overpicked and the rarer Targets 3 and 4 are underpicked (all $t(11) > 2.7$, all $p < 0.02$). Target 2 picking rate does not differ significantly from its prevalence rate ($t(11) = 0.87, p = 0.4$).

Why are observers favoring the more common item in this condition? This probably reflects the role of priming in hybrid foraging. As reported in Kristjansson et al. (2014) and Wolfe et al. (2016), selecting one target type biases the next selection toward the same target type. Since Target 1 is more common,

it tends to be found by chance more often, and the priming effect boosts its chances even further.

Uneven value, equal prevalence (Fig. 2b) In this condition, we can see the effect of value as observers tend to overpick the most valuable items and underpick the two least valuable items. (Here, targets are numbered inversely to value, 1—red—being the least valuable, and 4—green—being the most valuable.) This is most evident with the most valuable (16 points) and least valuable (2 points) targets. An ANOVA based on the first five selections again reveals a large effect of target type, $F(3, 33) = 9.26, p = .0001, \eta_G^2 = .23$. Again, there is no effect of the pick versus actual variable, $F(1, 11) = 0.10, p = .75, \eta_G^2 = 0.0$, because the effects go in different directions for different target types. This time, the most valuable are overpicked, the less valuable are underpicked, and, again, the interaction is significant, $F(3, 33) = 9.0, p = .0002, \eta_G^2 = .34$. In contrast to the previous condition, post hoc t tests show that Target 1 is *underpicked*, $t(11) = 3.1, p = .011$, as is Target 2, $t(11) = 3.0, p = .012$, and the most valuable target, Target 4 is *overpicked*, $t(11) = 3.4, p = .006$. The Target 3 picking rate does not differ significantly from its prevalence rate, $t(11) = 1.3, p = 0.2$.

Within each patch, this trend to choose the more valuable item wanes as participants progress because the valuable items become increasingly rare in that patch. Nevertheless, people still tend to overpick the most valuable and underpick the least valuable. As a result, people usually collect most of the valuable items in each patch but leave behind up to half of the least valuable items when they move to the next patch (see Fig. 5).

Uneven value, unequal prevalence (Fig. 2c) This condition represents the interaction of prevalence and value that occurs when observers can choose to pursue rare but valuable targets or more common but less valuable ones. In this condition, observers begin by overpicking the two most valuable items and underpicking the two least-valuable items, though the high prevalence of low-value targets means that they are the still most commonly chosen. The trend favoring valuable targets fades as observers progress through the patch because the already rare valuable items become even rarer and you cannot pick what is not there. Interestingly, as the valuable items run out, the observers tend to leave the patch, as will be discussed later. In this condition, especially, participants are willing to leave a large percentage of the less valuable items behind when they move on to a new patch (see Fig. 5). As before, an ANOVA based on the first five selections reveals a large effect of target type, $F(3, 33) = 47.6, p < .0001, \eta_G^2 = .23$. Again, there is no effect of the picked versus actual variable, $F(1, 11) = 1.7, p = .22, \eta_G^2 = 0.0$,

because the effects go in different directions for different target types. Again, the interaction is significant, $F(3, 33) = 4.6$, $p = .0085$, $\eta_G^2 = .20$. Post hoc t tests show that least valuable target, Target 1 is *underpicked*, $t(11) = 2.3$, $p = .046$. The more valuable Targets 3 and 4 are *overpicked*, both $t(11) > 2.2$, $p < .05$. The Target 2 picking rate does not differ significantly from its prevalence rate, $t(11) = 1.0$, $p = .32$. The smaller p values in this case reflect the greater variability between observers (also reflected in the error bars of Fig. 3). Some observers collect only the most valuable items and then move to the next patch. Others tend to start with the valuable items and then move to collect a considerable proportion of the less valuable items. Individual differences would be interesting to investigate in this task, but we are underpowered to say anything beyond noting that uneven-value, unequal-prevalence conditions produce the most variable data.

Again, overall observers remain fairly consistent in their value-favoring strategy across patches. This means that high-value targets are consistently collected at a rate higher than their prevalence in the display, while low-value targets are collected at a rate lower than their prevalence in the display. Nevertheless, because of the high prevalence of low-value targets, low-prevalence targets are the most common targets selected in absolute terms.

Fig. 4 a Average patch leaving times. b Number of patches viewed in each condition. Each symbol denotes one observer. Error bars show ± 1 SD. (Color figure online)

When do observers leave a patch?

Figure 4a shows the average leaving time for each observer in each of the three conditions. Figure 4b shows the number of patches viewed by each observer in each condition. Recall that the point values were assigned to roughly equate the amount of time. The average amount of time per condition was 16 minutes. There were no significant differences between the three conditions, $F(1.923, 21.15) = 1.589$, $p > .05$, df Greenhouse–Geiser corrected. Thus, the number of patches per condition varies inversely with time in each patch, and it is meaningful to compare across conditions because we are always dealing with number of patches visited in about 16 minutes.

Two trends are visible. The first is that observers appear to leave patches more quickly, and, consequently, they view more patches, when both value and prevalence differ across target types. The effect on patch leaving time is statistically marginal, $F(1.752, 19.28) = 3.033$, $p = .077$, df Greenhouse–Geiser corrected. The effect on patches viewed is statistically significant, $F(1.066, 11.72) = 12.68$, $p = .0036$, df Greenhouse–Geiser corrected. The second trend is that observers are more variable when both value and prevalence differ across targets. That is, everyone does more or less the same thing when the value of all targets is the same. When value varies, so does observer behavior. As noted before, individual differences would be interesting to study in future research.

What are observers leaving behind, when they leave a patch?

Figure 5 shows the average percentage of each target type that was left on-screen when the observer moved to a new patch. Recall that observers were under no obligation to collect all of the targets; their goal was to collect points as rapidly as possible.

Again, two trends are visible. The first trend is that the choice of target changes with condition. In the even value, unequal prevalence condition, observers collect the common targets and leave a larger percentage of the uncommon items behind. This is a version of a standard prevalence effect (Wolfe et al., 2005). The effect of target type is significant, $F(2.165, 23.81) = 12.99$, $p = .0001$, Greenhouse–Geiser corrected. When targets have different values, the pattern is very different. Now, the valuable items (T4) are most thoroughly collected. Again, the effects of target type is significant for both of the uneven value conditions, equal prevalence: $F(1.32, 14.58) = 7.85$, $p = .0094$; unequal prevalence: $F(1.734, 19.07) = 12.99$, $p = .0437$, Greenhouse–Geiser corrected. The other trend is that observers become highly variable in their behavior with regard to the low-value items. Some observers

Fig. 5 Proportion of targets left on screen in all three conditions. T1–T4 indicate the four target types. When unequal, the prevalence of a target is given in the green boxes. Pink boxes give point values. Each data point shows a single observer. Error bars show ± 1 SD. (Color figure online)

choose to collect virtually all of them, while others collect almost none. In a two-way ANOVA, the main effects of condition and target type are not significant, because these different patterns of response are reflected in a large interaction between the variables, $F(6, 66) = 10.6$, $p < .0001$, $\eta_G^2 = 0.17$. One way to examine this interaction is to compare the proportion left behind for Targets 1 and 4 in the different conditions. When Target 1 is much more prevalent than Target 4 and their values are equal, observers leave a greater proportion of the rare Target 4 on the screen, $t(11) = 6.3$, $p < .0001$. When Target 4 is much more valuable than Target 1 and their prevalences are equal, observers leave a greater proportion of low-value Target 1 on the screen, $t(11) = 3.4$, $p < .006$. When Target 1 is much more prevalent and Target 4 is much more valuable, observers leave a greater proportion of the

common but worthless Target 1 on the screen, $t(11) = 2.3$, $p < .04$. The size of this effect remains substantial in the average data. Observers leave 39% of Target 1 and just 18% of Target 4. However, as can be seen in Fig. 5c, observers differ widely in their interest in collecting Target 1.

Why are people leaving the current patch?

Obviously, observers are not, in general, choosing to collect every target on the screen. What rule are they using to leave the current patch? A candidate answer is given by the marginal value theorem (MVT; Charnov, 1976). MVT predicts that the forager should move to the next patch when the instantaneous rate of return in the current patch drops below the average rate of return in the task. Recall that the instantaneous rate of return is the rate at which an observer is collecting targets at a specific moment during foraging in a patch. It could be defined in various ways. For instance, one could compute this as a function of clicks in a patch. If the average RT for the sixth click in a patch is 1.5 seconds, then the rate is $1/1.5$, or 0.67 items per second. Suppose that only 80% of those clicks fell on actual targets, then the instantaneous rate of return would be 0.67×0.8 , or 0.53 targets/second. In this experiment, the value of the items needs to be considered as well. Obviously, the rate of return is greater in terms of points if you have been collecting high-value items as opposed to low-value items. We are interested in the instantaneous rate at the time of patch leaving. Accordingly, rather than computing a rate for the first, second, third click, and so forth, we average the number of points collected for each click *working backwards from the final click in the patch*. The rate is calculated by dividing the average points by the average RT for that click. The average time for those clicks is plotted on the x-axis.

The average rate of return is simply the total number of points divided by the total time for the block. This rate is reduced by the “travel time” between patches, when observers cannot collect any targets (averages 4–7 seconds in this experiment). Figure 6 shows the relationship of the instantaneous rate of return to the average rate of return. The thick, black lines are the average over 12 observers. The last point on each of those functions is the average of the rates for all observers for the final collection from a patch. The penultimate point represents the penultimate selection, and so forth. The x-axis gives the average time, measured from the appearance of the patch. Thin lines represent the instantaneous rates for each observer for the last 10 clicks in a patch. The dashed horizontal line represents the average rate of return, obtained by dividing total points collected by total time. Note that the point values are normalized by dividing by the expected point value of a click on a random target. In the even value conditions, each item is worth 4 points. In the uneven-value, equal-prevalence case, the expected value of a random target is (2

Fig. 6 Average, normalized rate of point collection as a function of time in patch. Thin colored lines represent individual observers. Thick lines represent the average of those observers. Dashed horizontal line shows average rate of return over the entire task. (Color figure online)

+ 4 + 8 + 16)/4 = 7.5. We could have made the expected value in the second case 4.0 by using point values of 1.1, 2.1, 4.3, and 8.5, but presumably this would not have changed the results. It is the relative values, not the absolute values of targets that are important. Normalization in Fig. 6 simply serves to place all the rates on the same scale.

Several aspects of Fig. 6 are of interest. First, individual observers vary widely in the rate with which they collect points, though the decline over time seems to fall on a roughly consistent function across observers. Second, for the even-value/unequal-prevalence and the uneven-value/equal-prevalence condition, the average results are roughly as predicted by MVT. Once the instantaneous rate falls to the average rate, the average observer leaves for the next patch. In the uneven-value/unequal-prevalence condition, however, observers, on average appear to leave the patch sooner than predicted by MVT. This point is examined in more detail in Fig. 7.

Figure 7 shows the normalized instantaneous rate (y -axis) plotted against normalized overall rate (x -axis) for each observer in each of the three conditions. As above, the normalization plots the different point values into a common scale. In a figure of this sort, MVT predicts that the final point (solid symbols) will fall just below the diagonal line of equality. A

point below the line shows the instantaneous rate falling below the average rate. The penultimate point (open symbols) should fall on or above the line. In general, there should be a close match between the average and instantaneous rates for these points at the moment of patch leaving. Thus, the data points should cluster near the main diagonal of the graph. In Fig. 7a, the predictions of MVT are well met. In Fig. 7b, most of the data fit MVT quite well, though a couple of observers' data points fall well above the line, suggesting that those observers left patches when it would have been profitable to continue collecting. This tendency is much more marked in Fig. 7c, where about half of the observers are leaving early. These are observers who collect a few high-value items and then quickly leave the patch. Returning to Fig. 5, these are the same observers who are leaving almost all of the lower value targets in the display when they move to the next patch.

Interestingly, the pattern of results looks different if we replot Fig. 7 using the click rate—how fast are observers collecting targets—rather than the point rate. This is shown in Fig. 8.

For the even-value, unequal-prevalence condition, the results are, by definition, identical to those in Fig. 6. For the other two conditions, plotting the results in terms of clicks

Fig. 7 Normalized instantaneous rate (y -axis) plotted against normalized overall rate (x -axis). Each data point represents one observer. Solid points plot the final instantaneous rate. Open symbols plot the rate for the penultimate selection

Fig. 8 Same as Fig. 6 but, the rates are normalized clicks per second. This is the rate at which observers collect items, regardless of their value

per second causes the data to cluster near the diagonal line of equality, as predicted by a variant of MVT in which they based their decisions on the rate of acquisition but not the value of what was acquired. One can imagine how this would happen. Suppose two observers began collecting high-value items. As the high-value items begin to run out and become hard to find, one observer switches to lower value items, the rate of return falls because the value falls but they can continue collecting items at a relatively high rate of items per second. The second observer has decided that only the high-value items are worth collecting. This observer continues to try to find those high-value items, but it takes longer and longer as they become scarce. Once it takes too long, this observer moves to a new patch rather than moving to a lower value target. Both of these observers are following a form of MVT, but over different sets of targets. Observer 1 quits when the rate of collection for all target types falls below the average rate. Observer 2 quits when the rate of collection of the first, high-value item falls below average.

General discussion

In this hybrid foraging task, there are two broad questions for the forager to answer: What is the next item I should be picking within a patch, and when should I move to the next patch? Within the patch, we can see three forces at work, shaping the observer's decision about what to pick next. These are illustrated in Fig. 9. The figure portrays a moment in a task with three target types: The painting set, the computer cable, and the stuffed duck. Each instance of the painting set is worth twice as many points as the other two target types.

The figure portrays the moment after the forager has collected the duck, shown circled in the array of items. What will be the identity of the next item to be collected? There are at least three forces at work. First, all else being equal, the observer will tend to go to a nearby object rather than one farther away. That would favor the cable (a). There are two parts to this bias: (1) It is less effort to move to a nearby item and (2) it is more likely that the observer will have already attended to

an item near the most recent selection than to a more remote item. The second force is priming. Consistent with previous work, in the present data, we see that observers are inclined to continue picking items of the type that they had been picking. This is not an overwhelming effect. If priming or a desire to avoid switch costs were paramount (Monsell, 2003), observers would pick all of one item, then all of another, and so forth. Though that did happen in the conjunction condition of Kristjansson et al. (2014), in the present experiment, it did not. Related to this point, and consistent with the work on prevalence effects (Horowitz, 2017), observers tend to pick more of the most common items than of the rarer items. Importantly, the proportion of common items that is picked is greater than the proportion of common items in the display. In Fig. 9, priming would favor picking another duck (b). The priming effect in foraging is described in more detail in Kristjansson et al. (2014) and Wolfe et al. (2016). The third force is value, the main topic of this article. In Fig. 9, value would favor picking the painting set (c).

In principle, it should be possible to predict the next selection with a considerable degree of precision, though that would require more information than is available in the present experiment. In particular, one would need to know the positions of every item at each moment. Recall that in this experiment, items are in continuous motion, and targets are

Fig. 9 A moment in a hybrid search task. The observer has just collected the duck in the circle. Which item will be selected next? (Color figure online)

removed from the display once collected. We did not preserve the information about the moment-by-moment state of the display that would be necessary to test precise predictions. Nevertheless, we can sketch how a precise model would work.

- A. The rate of collection is much slower than the rate at which items are identified or even the rate at which they are fixated. Items like the objects that are used here are probably identified at a rate of about 30 per second (Vickery, King, & Jiang, 2005; Wolfe, Alvarez, Rosenholtz, Kuzmova, & Sherman, 2011). Three or four of those items can be fixated each second. The rate of foraging is roughly one per second. Thus, the next item to be selected will come from a set of candidates. In the present experiment, a patch starts with 20–30% targets, so 20–30% of items attended at the start of the patch will be candidates for selection. As foraging progresses and the patch is depleted, the percentage will drop, so the size of the set of attended targets to collect will drop.
- B. The set of items attended will probably not be random. Introspectively, if you attend to one telephone in Fig. 9, all of the telephones seem to “light up” in a manner that suggests that your attention is being guided to the features of the phone (Wolfe & Horowitz, 2017). In the physiological literature, this would be called feature-based attention (Bichot & Schall, 2002; Maunsell & Treue, 2006; Treue, 2014). A precise model would need to make some assumptions about the trade-off between items selected because of their similarity to recently selected items and items selected because of their proximity to the current locus of attention and/or the most recently collected items. It is worth noting that the value of an item, independent of its features, probably does not guide attention (Rajsic, Perera, & Pratt, 2017).
- C. One would also need to model when, relative to the current collection, the observer is committed to the next collection. We can be quite sure that search for the next item begins well before the click that collects the current item. For example, if the search for the next item only starts after the click on the current target, then it would not matter if all the items change position and/or identity at the moment of that click. However, it does matter. The rate of foraging is much higher if the items are stable, rather than changing (Wolfe, Cain, Ehinger, & Drew, 2015).
- D. Data from experiments like those described here can be used to estimate the strength of proximity, priming, and value. The combination of those factors could be used to generate a probability map that should give a quite precise prediction about which item will be collected next. Returning to Fig. 9, if telephones were worth 20 points and the observer had just clicked on the phone in the upper right, the account, presented here, would have no

problem predicting that the next selection would be the phone down and to the left of the most recent collection. It would be most proximal, most valuable, and primed. The choice after the selection of the circled duck is less clear since the different factors are pitted against each other.

It is likely that there will be individual differences in the relative strength of proximity, priming, and value. These differences will influence the answer to the second broad question, raised at the start of the Discussion: How does value influence when observers move from one patch to the next? Suppose value is the strongest force for one forager. He will tend to be strongly inclined to collect all the high-value items first. As those items become scarcer, it will take longer and longer to find them. This will cause the instantaneous rate of target and point collection to fall and, in accordance with the MVT, the forager might depart for the next patch, leaving most or all of the lower value items in the current patch uncollected. This observer will have had a high rate of return that falls off relatively quickly.

In contrast, imagine a forager who is less moved by the value of the items. This observer will probably still favor valuable items over less valuable items, but as she picks targets, proximity and/or priming will cause some lower value items to be picked. As a result, her instantaneous rate of return will be lower (dimes and nickels will be mixed in with the quarters). However, because she is willing to pick more types of targets, there will be more targets to pick, and the instantaneous rate will fall off more slowly. This observer will stay longer in the current patch and, when she leaves, she will leave fewer items uncollected. This range of behavior can be inferred from Fig. 6, where the thin lines show the instantaneous rate for each observer. In Figs. 6b and 6c, where the value varies across target types, the family of functions fall on a decelerating curve. To the left are observers with high rates of return who quit after a relatively short time. On the right are those with lower rates who stay longer.

What is the best strategy? In this particular experiment, the goal was to get the required number of points as quickly as possible. As it happens, the observers who picked the valuable items and then left the patch finished the task faster than observers who picked a mix of targets. They also tend to pick at the fastest rate, making it hard to tell if their advantage came from picking fast or picking only valuable items. Indeed, they may have been able to pick fast because they essentially ignored the lower valued items. It will take further experiments to tease apart these factors. It seems unlikely that observers adjusted their selection rules dramatically in an attempt to optimize their performance. The time in each condition and the level of feedback seem inadequate. It is more likely that observers came into the experiment with implicit “settings” for value, priming, and proximity and these shaped their

behavior. It would be interesting to more directly influence those settings during the experiment. For instance, how would different observers respond to changes in the “travel time” between patches. MVT predicts that a longer travel time would produce slower patch leaving. However, it is not clear if that would be adequate to persuade an observer who was inclined to pick only high value items to collect lower value items. Going the other direction, if one target type was much more valuable than any other, would all observers adopt a strategy of leaving when those items became scarce?

Finally, it is interesting to speculate about how these forces might impact behavior in the world. For example, consider foraging in the supermarket. Some items are on sale, increasing their relative value. It would be unsurprising to find that those items are now collected at a higher rate. But how do the sale items influence collection of nonsale items? It would be interesting if differences in performance on a task like the one presented here were related to differences in behavior in the field. Would observers who left the patch when the high value items ran low, leave the store with only sale items in the basket?

Summary

Value modulates foraging behavior in human observers. As one would suspect, observers prefer more valuable targets. Interestingly, they seem to differ in the way in which that preference impacts their interest in less valuable targets. Some observers seem inclined to collect “only the best,” while others will collect less valuable items, perhaps in an effort to maximize their overall yield.

Acknowledgements This work was supported by NIH EY017001 and U.S. Army (NSRDEC) W911QY-16-2-0003.

References

- Bichot, N. P., & Schall, J. D. (2002). Priming in macaque frontal cortex during popout visual search: Feature-based facilitation and location-based inhibition of return. *Journal of Neuroscience*, 22(11), 4675–4685.
- Bond, A. B. (1981). Giving-up as a Poisson process: The departure decision of the green lacewing. *Animal Behaviour*, 29, 629–630.
- Bond, A. B., & Kamil, A. C. (2002). Visual predators select for crypticity and polymorphism in virtual prey. *Nature*, 415(6872), 609–613.
- Brady, T. F., Konkle, T., Alvarez, G. A., & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proceedings of the National Academy of Sciences of the United States of America*, 105(38), 14325–14329.
- Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10, 433–436.
- Cain, M. S., Vul, E., Clark, K., & Mitroff, S. R. (2012). A Bayesian optimal foraging model of human visual search. *Psychological Science*, 23(9), 1047–1054. <https://doi.org/10.1177/0956797612440460>
- Chan, L. K. H., & Hayward, W. G. (2012). Visual search. *WIREs*. Advance online publication. <https://doi.org/10.1002/wcs.1235>
- Chamov, E. L. (1976). Optimal foraging, the marginal value theorem. *Theoretical Population Biology*, 9(2), 129–136.
- Della Libera C., & Chelazzi L. (2006). Visual selective attention and the effects of monetary reward. *Psychological Science* 17, 222–227.
- Failing, M., & Theeuwes, J. (2017). Selection history: How reward modulates selectivity of visual attention. *Psychonomic Bulletin & Review*. <https://doi.org/10.3758/s13423-017-1380-y>
- Hickey, C., Chelazzi, L., & Theeuwes, J. (2010). Reward changes salience in human vision via the anterior cingulate. *Journal of Neuroscience*, 30, 11096–11103.
- Horowitz, T. S. (2017). Prevalence in visual search: From the clinic to the lab and back again. *Japanese Psychological Research*, 59(2), 65–108. <https://doi.org/10.1111/jpr.12153>
- Ishihara, I. (1980). Ishihara's Tests for Color-Blindness: Concise Edition. Tokyo: Kanehara & Co., LTD
- Kristjansson, A. (2006). Simultaneous priming along multiple feature dimensions in a visual search task. *Vision Research*, 46(16), 2554–2570.
- Kristjansson, Å., Johannesson, O. I., & Thornton, I. M. (2014). Common attentional constraints in visual foraging. *PLOS ONE*, 9(6), e100752. <https://doi.org/10.1371/journal.pone.0100752>
- Maljkovic, V., & Nakayama, K. (1994). Priming of popout: I. Role of features. *Memory & Cognition*, 22(6), 657–672.
- Maunsell, J. H., & Treue, S. (2006). Feature-based attention in visual cortex. *Trends in Neurosciences*, 29(6), 317–322. <https://doi.org/10.1016/j.tins.2006.04.001>
- Monsell, S. (2003). Task switching. *Trends in Cognitive Sciences*, 7(3), 134–140.
- Munneke, J., Hoppenbrouwers, S., & Theeuwes, J. (2015). Reward can modulate attentional capture, independent of top-down set. *Attention, Perception, & Psychophysics*, 77(8), 2540–2548. <https://doi.org/10.3758/s13414-015-0958-6>
- Navalpakkam, V., Koch, C., Rangel, A., & Perona, P. (2010). Optimal reward harvesting in complex perceptual environments. *Proceedings of the National Academy of Sciences of the United States of America*, 107, 5232–5237.
- Olivers, C. N. L., & Hickey, C. (2010). Priming resolves perceptual ambiguity in visual search: Evidence from behaviour and electrophysiology. *Vision Research*, 50(14), 1362–1371. <https://doi.org/10.1016/j.visres.2009.11.022>
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437–442.
- Pirolli, P. (2007). Information foraging theory. New York, NY: Oxford U Press.
- Pyke, G. H., Pulliam, H. R., & Chamov, E. L. (1977). Optimal foraging: A selective review of the theory and tests. *The Quarterly Review of Biology*, 52(2), 137–154.
- Rajsic, J., Perera, H., & Pratt, J. (2017). Learned value and object perception: Accelerated perception or biased decisions? *Attention, Perception, & Psychophysics*, 79(2), 603–613. <https://doi.org/10.3758/s13414-016-1242-0>
- Schneider, W., & Shiffrin, R. M. (1977). Controlled and automatic human information processing: I. Detection, search, and attention. *Psychological Review*, 84, 1–66.
- Serences, J. T. (2008). Value-based modulations in human visual cortex. *Neuron*, 60, 1169–1181.
- Stephens, D. W., & Krebs, J. R. (1986). Foraging theory. Princeton, NJ: Princeton University Press.
- Treue, S. (2014). Object- and feature-based attention: Monkey physiology. In A. C. Nobre & S. Kastner (Eds.), *Oxford handbook of attention* (pp. 573–600). New York, NY: Oxford University Press.
- Vickery, T. J., King, L.-W., & Jiang, Y. (2005). Setting up the target template in visual search. *Journal of Vision*, 5(1), 81–92.

- Wolfe, J. M. (2012). Saved by a log: How do humans perform hybrid visual and memory search? *Psychological Science*, 23(7), 698–703. <https://doi.org/10.1177/0956797612443968>
- Wolfe, J. M. (2013). When is it time to move to the next raspberry bush? Foraging rules in human visual search. *Journal of Vision*, 13(3), 10. <https://doi.org/10.1167/13.3.10>
- Wolfe, J. M. (2014a). Approaches to visual search: Feature integration theory and guided search. In A. C. Nobre & S. Kastner (Eds.), *Oxford handbook of attention* (pp. 11–55). New York, NY: Oxford University Press.
- Wolfe, J. M. (2014b). Visual Search. In E. R. A. J. F. Alan Kingstone (Ed.), *Handbook of attention*. Cambridge, MA: MIT Press.
- Wolfe, J. M., Aizenman, A. M., Boettcher, S. E., & Cain, M. S. (2016). Hybrid foraging search: Searching for multiple instances of multiple types of target. *Vision Research*, 119, 50–59.
- Wolfe, J. M., Alvarez, G. A., Rosenholtz, R., Kuzmova, Y. I., & Sherman, A. M. (2011). Visual search for arbitrary objects in real scenes. *Attention, Perception, & Psychophysics*, 73(6), 1650–1671. <https://doi.org/10.3758/s13414-011-0153-3>
- Wolfe, J. M., Cain, M., Ehinger, K., & Drew, T. (2015). Guided Search 5.0: Meeting the challenge of hybrid search and multiple-target foraging. *Journal of Vision*, 15(12), 1106–1106.
- Wolfe, J. M., & Horowitz, T. S. (2017). Five factors that guide attention in visual search. [Review article]. *Nature Human Behaviour*, 1, 0058. <https://doi.org/10.1038/s41562-017-0058>
- Wolfe, J. M., Horowitz, T. S., & Kenner, N. M. (2005). Rare targets are often missed in visual search. *Nature*, 435(7041), 439–440. <https://doi.org/10.1038/435439a>