The Distribution of People to Resources in a Networked Multi-player Environment

Robert L. Goldstone (rgoldsto@indiana.edu)
Benjamin C. Ashpole (bashpole@indiana.edu)
Cognitive Science Program
Indiana University
Bloomington, IN. 47405

Absract

We have developed an experimental platform that allows a large number of human participants to interact in real-time within a common virtual world. Within this environment, human participants foraged for resources distributed in two pools. In addition to varying the relative replenishment rates for the two resources (50-50, 65-35, 80-20), we manipulated whether participants could see each other and the entire food distribution, or had their vision restricted to food at their own location. Two empirical violations of optimal distributions of participants were found. First, there was a systematic underutilization of the more preponderant resource. Second, there were oscillations in the harvesting rate of the resources in a pool across time, revealed by a Fourier analysis with prominent power in the range of 50 seconds per cycle. These suboptimalities were more prevalent when participants had no knowledge of other participants or the complete food distribution. Individual participant knowledge thus affects the efficiency with which a population of participants harvests resources.

Introduction

A problem faced by all mobile organisms is how to search their environment for resources. Animals forage their environment for food, web-users surf the internet for desired data, and businesses mine the land for valuable minerals. When an individual organism forages their environment for resources, they typically employ a form of reinforcement learning to allocate their foraging time to the regions with the highest utility in terms of providing resources (Ballard 1997). As the organism travels in its environment, it gathers information about resource distributions and uses this information to inform its subsequent movements. When an organism forages in an environment that consists, in part, of other organisms that are also foraging, then additional complexities arise. The resources available to an organism are affected not just by the foraging behavior the organism itself, but also by the simultaneous foraging behavior of all of the other organisms. The optimal resource foraging strategy for an organism is no longer a simple function of the distribution of resources and movement costs, but rather is also a function of the strategies adopted by other organisms. Even if the resources are replenished with a constant rate, the optimal foraging strategy for an organism may fluctuate depending on the other organisms' behavior.

We have developed an experimental platform that allows many human participants to interact in real-time within a common virtual world. Inside this world, two resource pools were created, and we recorded the moment-by-moment exploitation of these resources by each human agent. The research questions motivating the current study are: "How do resource foraging strategies unfold with time?," "Are there systematic suboptimalities in resource foraging?," and "How are foraging strategies affected by the distribution of resources and the agents' knowledge of the environment?"

Foraging in Groups

Determining how to allocate time and energy to resources is a deep issue that also has practical importance in biology, economics, psychology, and computer science. Biologists have long entertained the hypothesis that animals forage for resources in a near-optimal manner, given the distribution and replenishment rate of the resources, the animals' resource demands, and the energy expenditures required to harvest the resources (Krebs & Davies, 1978). Biologists studying individual foraging behavior have found many situations where resource patches are visited by animals with efficiency (Pleasants, 1989). For example, hummingbirds have been shown to sample the rate of return of flowers in a region, then forage one flower until the rate of return is below the average for all flowers, and then forage another flower with greater-than-average return (Pyke, 1978).

These analyses from biology are based on a single individual harvesting resources without competition from other agents. This assumption is unrealistic in a world where organisms typically congregate in populations. Biologists have also explored the allocation of a population of organisms to resources. One model in biology for studying the foraging behavior of populations rather than individuals is Ideal Free Distribution (Fretwell & Lucas, 1972). According to this model, animals are free to move between resource patches and have correct ("ideal") knowledge of the rate of food occurrence at each patch. The animals distribute themselves to patches where the gained resources will be maximized. The patch that maximizes resources will depend upon the utilization of resources by all agents. Consistent with this model, groups of animals often distribute themselves in a nearly optimal manner, with their distribution matching the distribution of resources. For example, Godin and Keeleyside (1984) distributed edible larvae to two ends of a tank filled with cichlid fish. The food was distributed in ratios of 1:1, 2:1, or 5:1. The cichlids quickly distributed themselves in rough accord with

the relative rates of the food distribution before many of the fish had even acquired a single larva and before most fish had acquired larvae from both ends. Similarly, mallard ducks distribute themselves in accord with the rate or amount of food thrown at two pond locations (Harper, 1982). If one patch produces two times the amount of resource as another patch, there will be about two times as many animals at the larger relative to smaller resource.

The current research explores the foraging behavior of groups of humans. In general accord with an Ideal Free Distribution model, groups of fish, insects, and birds have been shown to efficiently distribute themselves. Are groups of people as rational as these animals? We were interested in manipulating the relative outputs of the different resource patches and the knowledge possessed by the agents. With respect to the first manipulation, the central issue is whether systematic inefficiencies in a population of agents arise as a function of the distribution of resources. Although Godin and Keenleyside found that cichlids approximately distribute themselves in accord with the food resources, they also found small but systematic deviations. For the 1:1 and 2:1 ratios of resources, the fish distributed themselves with approximately 1:1 and 2:1 ratios respectively. However, for the 5:1 ratio, the fish distributed with a ratio of approximately 3:1. This pattern of agents distributing themselves to resources in a less extreme (compared to a baseline of uniform distribution) manner than the resources themselves are distributed is called "undermatching." If undermatching occurs, then an adviser (perhaps a cichlid efficiency consultant) would recommend that a fish partaking of the resource with the lower output rate could increase its resource intake by moving to the resource with the higher output rate.

The few experiments that have examined group foraging behavior with humans have also found undermatching (Kraft & Baum, 2001; Madden, Peden, & Yamaguchi, 2002; Sokolowski, Tonneau, & Freixa i Baque, 1999). In some of these experiments (Kraft & Baum, 2001; Madden et al., 2002), each participant in a group holds up a card indicating one of two colors, and then one color is probabilistically selected as the rewarded card. The participants who are holding up the rewarded color share the rewards. In another paradigm (Madden et al., 2002), participants move between two designated areas and share rewards that are probabilistically deposited in these areas. In both paradigms, greater average amounts of reward are obtained by participants at the more productive resource, indicating undermatching.

Our experiment extends the previous studies of group foraging in humans in a few directions. First, we have developed a computer-based platform for the foraging experiment that allows us to manipulate experimental variables that would be difficult to manipulate in a more naturalistic environment. Second, we collect second-by-second data on the number of resources and participants at different resources, which allows us to explore variation in resource use with high temporal resolution. Third, although our environment is virtual, it is naturalistic in one important respect – resources are distributed in a continuous spatial environment rather than at two discrete locations. Fourth, we do not designate or identify the resource alternatives to participants. As in many natural situations (Kennedy & Gray, 1993), the organisms must discover the number and locations of resource patches themselves.

The other experimental variable that we manipulate is agents' knowledge of their environment and other agents. In Godin and Keenleyside's experiment with cichlids, every cichlid could see the other cichlids as well as the larvae resources at both ends of the tank. Gallistel (1990) argues that this information is important for the cichlids to distribute themselves rapidly in accord with the resource distribution. They are learning about the resource distributions by observing events that do not directly involve themselves. However, in standard reinforcement learning situations, an agent only has access to the outcomes of their own actions. They do not have access to the values of options not selected. Both situations occur naturally, and it is possible that the ability of a group to efficiently distribute itself to resources depends on the knowledge at each agent's disposal. It is plausible to hypothesize that as the information available to agents increases, the efficiency with which they can allocate their energy to resources increases, although there are paradoxical cases where more information seems to lead to worse decisions (Gigerenzer, Todd, & The ABC Research Group, 1999).

Experiment in Group Foraging

We have developed a software system that records the instant-by-instant actions of individuals foraging for resources in a shared environment. Two resource pools were created with different rates of replenishment. The participants' task was to obtain as many resource tokens as possible during an experiment. An agent obtained a token by being the first to move on top of it. In addition to varying the relative replenishment rate for the two resources (50-50, 65-35, or 80-20), we manipulated whether agents could see each other and the entire food distribution, or had their vision restricted to food at their own location. We were interested in analyzing the resulting data for dynamics and sub-optimalities in the allocation of individuals to resources.

Experimental Methods

One-hundred and sixty-six undergraduate students from Indiana University served as participants in order to fulfill a course requirement. The students were run in 8 groups with 21, 20, 23, 19, 28, 12, 25, and 18 participants. Each student within a group was assigned to a PC computer in a large computer-based classroom with 40 computers. The experimenter controlled the experiment from another computer in the same room. The participants' computers were registered with the experimenter's computer using our developed software for sending messages over the internet to a set of networked computers.

Participants were told that they were being asked to participate in an experiment on group behavior. They were instructed to try to pick up as many "food" resources as possible by moving their icons' position on top of food locations. Participants were told that the food would occur in clumps, and if they learned where the productive clumps were that they could harvest more food. Participants were also told that there would be a lottery at the end of the experiment to win $10, and every piece of food that they collected during the experiment would be worth one lottery ticket. In this manner, participants were motivated to collect as many pieces of food as possible, even if at some point in the experiment they felt that it was improbable that they could collect more food than any other participant. Participants were told to only look at their own computer screen and not to talk with each other at any time.

Participants within a group co-existed in a virtual environment made up of replenishing resource pools and other human-controlled agents. The environment consisted of an 80 X 80 grid of squares. Participants controlled their position within this world by moving up, down, left, and right using the four arrow keys on their computers' keyboard. A participant could not walk off one side of the grid and reappear on the other. Each participant was represented by a yellow dot. In the "Visible" condition, all of the other participants' locations were represented by blue dots, and available food resources were represented by green dots. In the "Invisible" condition, each participant saw only their own position on the screen and any food gathered by that participant in the last two seconds. After this time interval, these consumed food pieces disappeared. In both conditions, food was gathered when a participant's position coincided with a piece of food.

Every experiment was divided into six 5-minute sessions. These six games consisted of all combinations of the two levels of knowledge (Visible versus Invisible) and the three levels of resource distribution (50/50, 65/35, 80/20). For each of the three distribution conditions, two resource pools were constructed, with center locations at reflections and rotations of the set of coordinates {40,15} and {15, 65}. A different reflection and rotation was used for each of the six conditions, with the result that the resource centers were approximately equally likely to be in each of eight possible locations, and the two centers within one session always had the same distance from one another. Two opposite orders of the six games were randomly assigned to the eight separate groups of participants. Each order began with a visible condition and then alternated between visible and invisible conditions.

The rate of distribution of food was based on the number of participants, with one piece of food delivered every 4/N seconds, where N is the number of participants. This yields an average of one food piece per participant per four seconds. When a piece of food was delivered, it was assigned to a pool probabilistically based upon the distribution rate. For example, in the 80/20 condition, the food would occur in the more plentiful pool 80% of the time, and in the less plentiful pool 20% of the time. The location of the food within the pool followed a Gaussian distribution with a mean at the center of the pool and a variance of 5 horizontal and vertical positions. Thus, the probability of food occurring in a given location was inversely proportional to the distance between the location and pool's center. Since multiple agents could occupy the same location without colliding, any food placed on such a location would be randomly assigned to one of the agents at that location.

Each of the six sessions lasted 5 minutes. Data were recorded every two seconds that included the positions of all agents, the number of food pieces collected by each agent, and the locations of uncollected food pieces. After all six sessions were completed, a winning participant was selected to receive $10 by an automatic lottery. The probability of a participant winning the lottery was equal to the number of food pieces they collected divided by the total number of food pieces collected by all participants.

Results

As a preliminary analysis of the distribution of agents across resource pools, Figure 1 shows the frequency with which each of the 80 X 80 grid cells was visited by participants broken down by the six experimental conditions. The
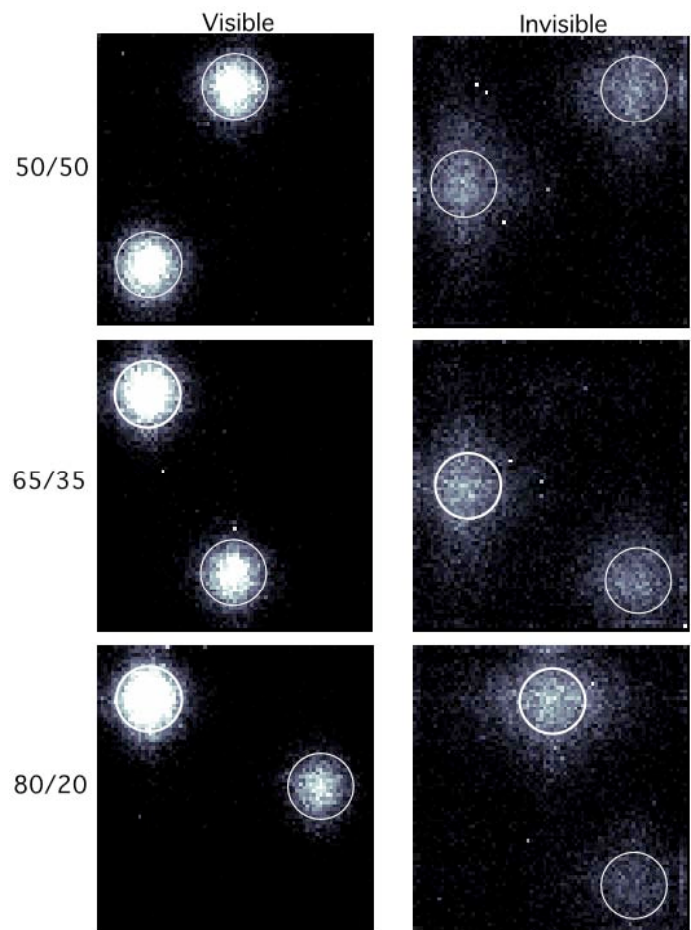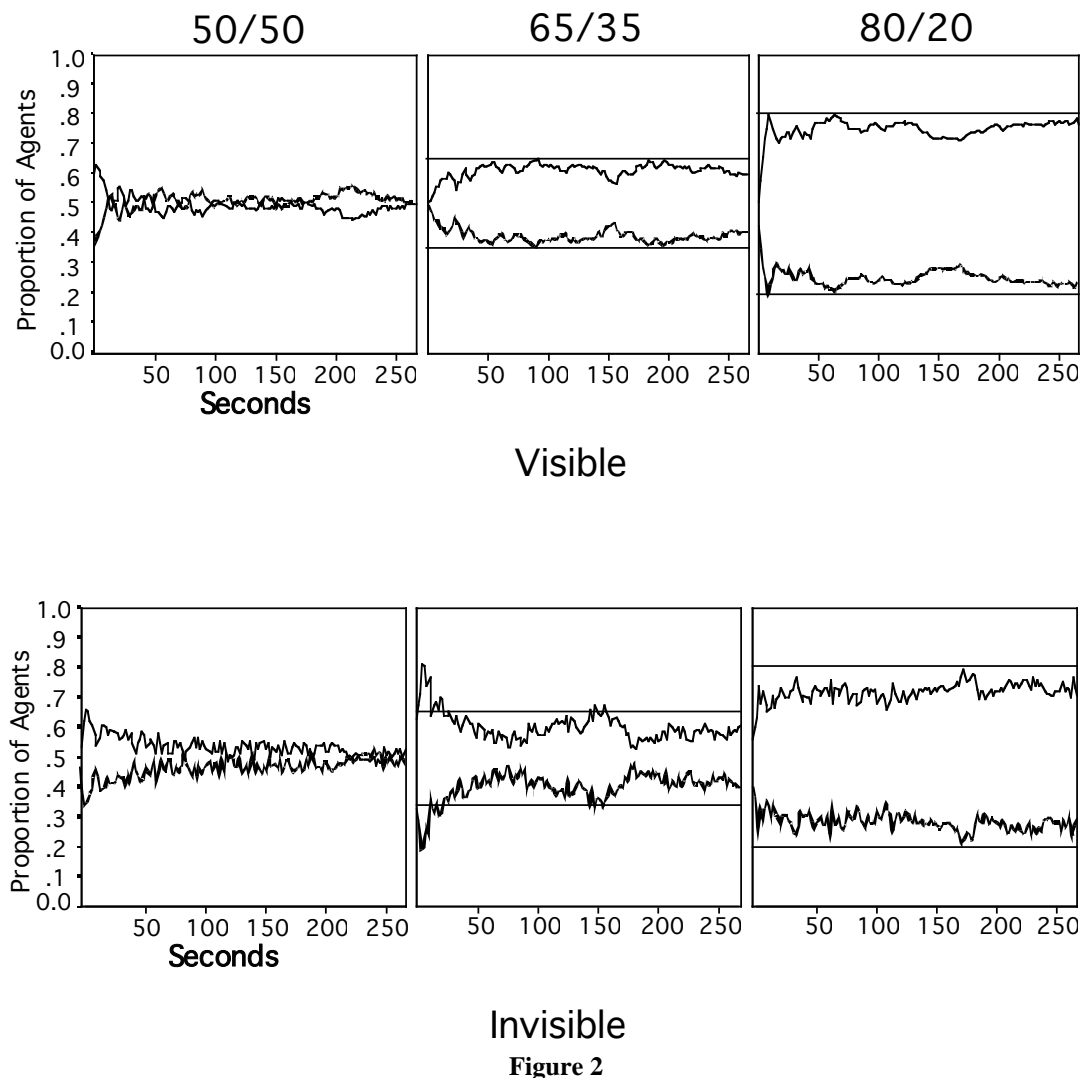

Figure 1

464

brightness of a cell increases proportionally with the number of times the cell was visited. The few isolated white specks can be attributed to participants who decided not to move for extended periods of time. In Figure 1, the thick and thin circles show two standard deviations of the food distribution for the more and less plentiful resources, respectively. An inspection of this figure indicates that agents spend the majority of their time within relatively small regions centered on the two resource pools. The concentration of agents in pools' centers is greater for visible than invisible conditions, and is greater for the more plentiful pool. For the invisible conditions, there is substantial diffusion of travel outside of one standard deviation of the pools' centers. A Cochran's test for homogeneity of variances revealed significantly greater variability for the invisible than visible condition (p<.01), indicating greater scatter of agents' locations in the invisible condition. The agents approximately distributed themselves in a Gaussian form, with the exception of a small second hump in the frequency distribution in the invisible condition. The cause of this hump is that cells near the edges of the 80 X 80 grid that

were close to pools were frequented somewhat more often than cells closer to the pool's center.

The dynamics of the distribution of agents to resources is shown in Figure 2, broken down by the six conditions. In this figure, the proportion of agents in the two pools is plotted over time within a session. Horizontal lines indicate the proportions that would match the distribution of food. An agent was counted as residing in a pool if he/she was within 5 food distribution standard deviations of the pool's center. This created circular pools that were as large as possible without overlapping. Agents who were not in either pool were excluded from Figure 2, and the total number of agents was normalized to exclude these agents. Figure 2 shows that agents roughly match the food distribution probabilities and that asymptotic levels of matching are found within 40 seconds even for the invisible condition. Although fast adaptation takes place, the asymptotic distribution of agents systematically undermatches the optimal probabilities. For the 65/35 distribution the 65% pool attracts an average of 60.6% of the agents in the 50-300 second interval, a value that is
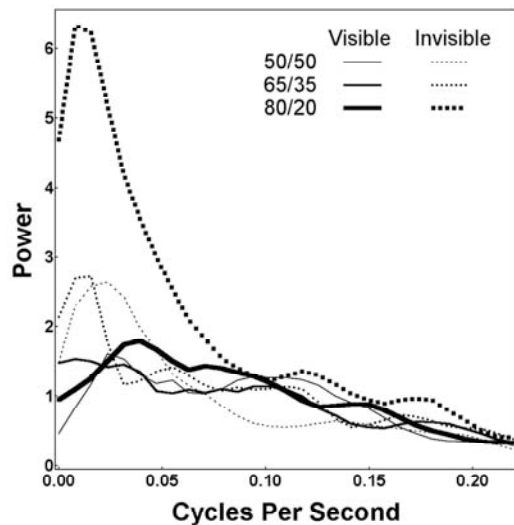

Figure 2

Figure 3

significantly different from 65%, one-sample T-test, t(7)=3.9, p<.01. Similarly, for the 80/20 distribution, the 80% pool attracts only 73.5% of the agents, t(7)=4.3, p<.01. For the 65/35 distribution, the asymptotic percentage of agents in the 65% pool in the visible condition (61.3%) was greater than in the invisible condition (60.0%), paired T-test t(7)=2.4, p < .05. Likewise, for the 80/20 distribution, the asymptotic percentage of agents in the 80% pool in the visible condition (74.8%) was greater than in the invisible condition (72.2%), t(7)=2.9, p < .05. Another trend, apparent in Figure 2, is that the proportions of agents in a given pool vary more sporadically with time for the invisible than visible conditions. This is because agents more often move themselves outside of a designated pool in the invisible condition. The percentage of agents falling outside of either pools during the interval 50-300 seconds were 1.2% and 13.4% for visible and invisible conditions respectively, paired t(7)=8.4, p<.01.

A final analysis of interest explores the possibility of periodic fluctuations in resource use. Informal experimental observations suggested the occurrence of waves of overuse and underuse of pools. Participants seemed to heavily congregate at a pool for a period of time, and then became frustrated with the difficulty of collecting food in the pool (due to the large population in the pool), precipitating an emigration from this pool to the other pool. If a relatively large subpopulation within a pool decides at roughly the same time to migrate from one pool to another, then cyclic waves of population change may emerge. This was tested by applying a Fourier transformation of the data shown in Figure 2. Fourier transformations translate a time-varying signal into a set of sinusoidal components. Each sinusoidal component is characterized by a phase (where it crosses the Y-intercept) and a frequency. For our purposes, the desired output is a frequency plot of the amount of power at different frequencies. Large power at a particular frequency indicates a strong periodic response.

Any periodic waves of population change that occur in the experiment would be masked in Figure 2 because the graphs average over 8 different groups of participants. If each group showed periodic changes that occurred at different phases, then the averaged data would likely show no periodic activity. Accordingly, to produce the frequency plot showed in Figure 3, we conducted four steps of analysis. First, we derived a data vector of the proportion of agents in the more plentiful pool across a five-minute session for each of the 8 groups within each of the 6 conditions. Second, we detrended each vector by removing the best straight-line fit from it. If we had not done this, the resulting frequency plot would exhibit inappropriately high power in low frequencies, reflecting slow trends in population growth or decline. Third, we applied a digital Fourier transformation (Fast-Fourier Transform) on each detrended vector. Fourth, we created the frequency plots in Figure 3 by averaging the frequency plots for the 8 groups within a single condition.

The resulting frequency plots show significantly greater power in the low frequency spectra for invisible than visible conditions. In particular, the average power for frequencies up to .05 cycles is 3.4 and 1.1 for invisible and visible conditions, paired T-test t(7)=4.1, p<.01. The power in lower frequencies is particularly high for the invisible condition with an 80/20 distribution. For all three invisible conditions, the peak power is at approximately .02 cycles/second. This means that the agents tend to have waves of relative dense crowding at one pool that repeat about once every 50 seconds. This 50 second period includes both the time to migrate from the first pool to the second pool and to return to the first pool. A pronounced power peak at lower frequencies is absent for the visible condition. One account for the difference in the two visibility conditions is that in the visible condition, each agent can see the whether other agents are closer than themselves to underexploited resource pools. The temptation to leave a dissatisfying pool for a potentially more lucrative pool would be tempered by the awareness that other agents are already heading away from the dissatisfying pool and toward the lucrative pool. However, in the invisible condition, agents may become dissatisfied with a pool populated with many other agents, but as they leave the pool they would not be aware that other agents are also leaving. It is less clear why periodic population waves should be greatest for the most lop-sided, 80/20 distribution, but one speculation is that the slowly replenishing 20% pool has the least power to attract a stable population of agents. A few agents can come into the 20% pool and quickly exhaust all of the fallen food resources. Then, the slow replenishment rate gives all of the agents in the pool little incentive to stay, and they consequently move to the 80% pool, until eventually the 20% pool becomes attractive again because of its low population density and cache of accumulated food.

Conclusions

The results of the present experiment indicate three systematic inefficiencies in the distribution of human participants to resources over time. First, participants exhibited undermatching in the sense that there were too many participants in the less plentiful resource and too few participants in the more plentiful resource. This undermatching is implied by comparing the distribution of agents to the distribution of resources, and is also directly indicated by the lower rate of food intake for agents in less plentiful compared to more plentiful resource. If this result proves to be general, then advice could be given to participants in similar situations to increase their use of relatively plentiful resources despite the possibly greater population density at those resources. Second, systematic cycles of population change are apparent whereby the migration of people from one pool to another is roughly synchronized. The problem with these synchronized population waves is that competitive crowding results in decreased food intake for those participants moving in the wave. Third, participants were more scattered than were the food resources. Both participants and food were distributed in a roughly Gaussian form, but the positional variance associated with participants was higher.

All three of these systematic inefficiencies were more pronounced for invisible than visible conditions. In fact, of the three inefficiencies described above, the only one that was appreciably large in the visible condition was the systematic undermatching. The influence of visibility suggests that an individual's knowledge of the moment-by-moment state of the environment and other agents can allow the group as a whole to avoid inefficient waves of resource under- and over-exploitation.

A fruitful extension to the current work would be the development of an agent-base computer simulation that models the large quantity of spatial-temporal population information amassed in the experiment. Such a model would need to incorporate a distinction between agents with and without vision. Blind agents might resemble standard reinforcement learning devices if supplemented by biases that undermatch distributions. Incorporating agents with vision is more challenging, raising important issues in generating expectations and planning. However, the current empirical work suggests that developing these more sophisticated agents with knowledge is worth the trouble. Knowledge of food distributions allows an agent to more effectively match those distributions, whereas knowledge of other agents allows an agent to more effectively decouple their responses from others.

Acknowledgments

References

Ballard, D.H. 1997, An introduction to natural computation, MIT Press (Cambridge).

Fretwell, S. D. & Lucas, H. J. (1970). Ideal free distribution. Acta Biotheory 19, 16-21.

Gallistel, C.R., 1990, The organization of learning, MIT Press (Cambridge).

Gigerenzer, G., Todd, P.M., & the ABC Research Group, 1999, Simple heuristics that make us smart, Oxford University Press (Oxford).

Godin, M.J., & Keenleyside, M.H.A., 1984, Foraging on patchily distributed prey by a cichlid fish (Teleosti, Chichlideae): A test of the ideal free distribution theory, Animal Behavior, 32, 120.

Harper, D.G.C., 1982), Competitive foraging in mallards: Ideal free ducks, Animal Behavior, 30, 575.

Kraft, J. R., & Baum, W. M. (2001). Group choice: The ideal free distribution of human social behavior. Journal of the Experimental Analysis of Behavior, 76, 21-42.

Madden, G. J., Peden, B. F., & Tamaguchi, T. (2002). Human group choice: Discrete-trial and free-operant tests of idea free distribution. Journal of the Experimental Analaysis of Behavior, 78, 1-15.

Krebs, J.R., & Davies, N.B., 1978, Behavioural Ecology: An Evolutionary Approach. Blackwell (Oxford).

Pleasants, J.M. 1989. Optimal foraging by nectarivores: a test of the marginal-value theorem. American Naturalist, 134, 51.

Pyke, G.H. 1978. Optimal foraging in hummingbirds: testing the marginal value theorem. American Zoologist, 18, 739.

Sokolowski, M. B. C., Tonneau, F., & Freixa I Baque, E. F. (1999). The ideal free distribution in humans: An experimental test. Psychonomic Bulletin & Review, 6, 157-161.