

Explaining Resource Undermatching with Agent-Based Models

Michael E. Roberts (Roberts@m@Indiana.Edu)

Department of Psychology, 1101 East Tenth Street
Bloomington, IN 47405 USA

Robert L. Goldstone (Rgoldsto@Indiana.Edu)

Department of Psychology, 1101 East Tenth Street
Bloomington, IN 47405 USA

Abstract

We propose two agent-based models of group foraging for two perceptual conditions. These models exhibit complex group-level behavior using only a simple rule set with a homogeneous group of agents. The models are shown to replicate results from Goldstone and Ashpole (2004), and we describe a series of simulations that test the sources of the resource undermatching often found in group foraging experiments. After testing the effects of travel costs, the number of agents, and uniform variance food distributions, we conclude that many group foraging studies have overlooked the interplay of spatial constraints with food rates in causing undermatching.

Introduction

The ideal free distribution (IFD) model (Fretwell & Lucas, 1970) predicts that a group of foragers will distribute to resource patches in proportion to the relative resources available at each patch. In an environment where one resource pool holds 80% of the resources while a second pool holds the remaining 20% of resources, the IFD predicts that a group of foragers will optimally distribute themselves to the resource pools, with 80% of the foragers in the first pool and 20% in the second pool. In predicting this optimal distribution, the model assumes the foragers have both freedom of movement and ‘perfect’ of the locations, amounts, and appearance rates of all the resources, thereby allowing the foragers to correctly value the resources and optimally distribute as a group. However, in experiments with two patch options, many studies report systematic undermatching in which fewer than expected foragers attend the more profitable patch while more than expected foragers attend the less profitable patch. Godin and Keenleyside (1984) found undermatching in cichlid fish, with the largest degree of undermatching occurring in the condition with the largest rate difference between resource pools. Gillis and Kramer (1987) obtained similar results for large populations (120 and 240) of zebrafish distributing between three resource pools, although they did not find significant deviations for small populations (30 and 60).

Whereas Godin and Keenleyside concluded that differential competitive abilities had broken the IFD assumption of equal payoff for each animal at a pool, Gillis and Kramer concluded that high population density broke the IFD assumption of perfect information. Despite the

difference in emphasis, both studies seem to agree that perceptual limitations – whether due to individual ability or occlusion from high density – led to decreased available information about the pools and subsequently undermatching. However, other studies (Harper, 1982; Kennedy & Gray, 1993) cite competitive differences among organisms and travel costs between locations as critical factors in producing undermatching. Baum and Kraft (1998) consistently found undermatching in a group of 30 pigeons distributing between two resource patches, although the degree of undermatching varied as a factor of resource distributions and travel constraints, and they were unable to detect whether competitive abilities varied. In a meta-analysis of undermatching in animal foraging studies, Kennedy and Gray (1993) conclude that information regarding the “relative and absolute resource availability, number of animals, perceptual abilities of animals, competitive interactions, competitive abilities of animals, and the effects of travel between sites” may all lead to undermatching and violate the IFD.

In the present article, we describe two agent-based models that capture several phenomena of group foraging across perceptual and resource conditions. These models show that complex group-level behaviors can emerge from simple agent behaviors. Furthermore, the models allow us to test changes in dynamic group behavior as a result of environmental manipulations, and give rise to a novel explanation of undermatching in group foraging.

Agent-Based Models

Goldstone and Ashpole (2004) recently examined dynamic group foraging behavior among humans by using an experimental networked Java platform to create a common virtual world (an 80 x 80 grid) across computers. Participants sat at their respective computer monitors and foraged for resources in real time by using the computers’ arrow keys to move up, down, left, and right in order to step on a food pellet and thereby consume it. We will briefly describe the experimental manipulations because the Goldstone and Ashpole foraging environment and data serve as the basis for our subsequent agent models.

In the Goldstone and Ashpole foraging experiments, participants engaged in 6 five-minute sessions, consisting of all permutations of two perceptual conditions and three

resource conditions, and all participants experienced the same conditions in a given session. In the “visible” perceptual condition, a participant could see himself or herself as a yellow dot in the virtual world, and other participants were visible as blue dots while available food pellets were represented as green dots. In the “invisible” perceptual condition, a participant could see himself or herself as a yellow dot in the virtual world, but no other participants or food were visible in the world. A new food pellet was dropped in one of two resource pools every $4/N$ seconds (where N is the number of participants), and there were three distribution conditions that probabilistically determined which pool received the pellet: 50/50, 65/35, and 80/20. For example, in the 65/35 distribution condition, 65% of food arrived at one pool while 35% arrived at the other pool. At each pool, new pellets were dropped according to a Gaussian distribution with a mean at the pool’s center and a variance of 5 units horizontally and vertically. Food release was constrained so that only one pellet could occupy a cell at a given time, and resource pool locations changed from session to session. In the invisible perceptual condition, a pellet appeared on the screen for two seconds for the participant who stepped on it, and therefore, participants could gradually ascertain the locations of the resources by exploring the world and occasionally obtaining pellets.

Goldstone and Ashpole discuss three results: undermatching, systematic cycles of population migration between pools, and high variance of participants’ locations relative to food locations. Although our agent-based models capture all of these phenomena using the same parameter values, we are most concerned with the undermatching for the present article. Goldstone and Ashpole found significant undermatching at the 80/20 and 65/35 distribution levels, with more pronounced undermatching in the more extreme 80/20 conditions. Since individuals do not directly interfere with each other, competitive interference can be removed as a possible explanation of the phenomenon. Furthermore, since undermatching occurs even in the visible conditions, where ‘perfect’ knowledge is available, straightforward accounts in terms of perceptual constraints can be removed as explanations. Our agent-based models seek to reproduce these results while testing the other undermatching alternatives – different competitive abilities, travel costs, and the number of foragers – from Kennedy and Gray’s (1993) meta-analysis.

Visible Model

An agent-based model should have minimal representational requirements while also being sufficiently robust to handle changes in the empirical environment, such as a sudden increase in food rates or number of participants. Gallistel (1990) has argued that a strong set of internal representations are necessary in any reasonable treatment of the IFD, but Magnuson (1998) proposed a relatively successful simple foraging model that utilized two rules: move to the nearest piece of food, and stay at a piece of

food for a number of time steps proportional to the magnitude of the food. In a favorable environment, these rules lead to resource matching; however, such ‘favorable’ environments are very constrained and very rare. Our visible model offers a compromise between the Gallistel and Magnuson approaches. Whereas Magnuson’s goal was simply to demonstrate the feasibility of group matching with simple rules, we propose an enhanced rule set in hopes of building a bridge between low-level agent behaviors and a collection of high-level group phenomena.

In our model, each agent is randomly assigned a location on the 80 x 80 grid-world at the beginning of the experiment. As in Goldstone and Ashpole (2004), food is dropped at a rate of $4/N$ seconds, and for all simulations reported in this article, N is set to 20 agents, corresponding to the average number of participants in the Goldstone and Ashpole experiment. Agents move every 100 milliseconds, and a movement consists of choosing an available food pellet and moving one grid unit towards the chosen pellet. The interesting aspects of the model lie in the parameters for action selection. During a time step, each agent judges the value of every available piece of food relative to the agent. Four factors determine a pellet’s worth: 1. Euclidean distance from the agent to the pellet, 2. Euclidean distance from the pellet to the agent’s sustained goal, 3. Food density in the local area of the pellet, and 4. Agent density in the local area of the pellet. Each factor is multiplied by a free parameter, yielding Equation 1 on the following page.

Individual factors can have varying importance depending on the current environment, but as a rule set, they produce robust, complex behavior. In particular, Euclidean distance captures the notion that animals will tend to approach the food nearest to them, since it is easiest to obtain and requires the least energy expenditure. The sustained goal factor (goal bias) makes an agent more likely to continue towards the previous goal pellet, thereby promoting inertia and countering a tendency for behavioral synchrony and generally unrealistic switching behavior found in Magnuson’s agents. The food density factor addresses another shortcoming of the Magnuson model. Magnuson’s agents were unable to differentiate between two resource pools with vastly different resource densities, despite the improved reward rate that can be obtained by such a distinction. In our model, if there are two unequally rewarding resource pools and an equal number of agents at each pool, food density will bias agents to leave the less plentiful pool in favor of the more plentiful pool. Finally, agent density is similar to food density, but it captures the idea that agents do indeed know the locations of other foragers in the visible condition, and they may want to avoid regions with high densities of competitors. In contrast to the food density factor, if two resource pools have equal available resources but different numbers of agents, agents in the dense pool should be encouraged to migrate to the other pool.

In addition to these factors, the model also incorporates the softmax action selection algorithm (Equation 2) to

Equation 1: $Pellet(i, j) = (P_1 * Distance) + (P_2 * GoalBias) + (P_3 * FoodDensity) - (P_4 * AgentDensity)$

$$\text{Equation 2: } Pr(i, j) = \frac{e^{Value(i, j)/K}}{\sum_x \sum_y e^{Value(x, y)/K}}$$

Equation 3: $Cell(i, j, t) = Cell(i, j, t-1) + (P_1 * Reward) + (P_2 * GoalBias) - (P_3 * Penalty)$

Box 1: Equations

probabilistically choose a pellet, given each pellet's relative value to the agent calculated in Equation 1. Although we have tried to make the expression generic to also accommodate the invisible model, we should point out that the denominator in Equation 2 only sums over pellets currently on the gridworld. In the softmax equation, K is a constant that shifts action selection to be more exploratory or more exploitative. When K is large, even highly valued pellets become small values and thereby indistinguishable, leading to uniform random movement choices. When K is small, differences between the pellet values are exaggerated by exponentiation, and the probability of choosing the highest valued pellet approaches 1 as K approaches 0. Together, the Euclidean measures and softmax action selection create a degree of agent independence while maintaining the essential dependence on resource and agent distributions for choosing where to move.

Invisible Model

For the invisible condition of Goldstone and Ashpole (2004), participants only saw a food pellet if they happened to step on it, so the perceptual factors from the visible model do not apply. Instead, foraging performance seems to rely on some form of memory, leading foragers to continue exploring regions where they have been rewarded in the past. Harley (1981) proposed a 'relative payoff sum' rule in which an agent's probability of choosing an option is equal to the proportion of past rewards from that option relative to the sum of past rewards from all options. Harley also included prior expectations and a memory decay factor in his model, and he successfully conducted computer simulations showing that a group could achieve IFD with the rule. Regelmann (1984) subsequently extended Harley's model to allow for competitive differences between individuals.

Our invisible model is somewhat similar to Harley and Regelmann's models, but in addition to component differences, it has several key differences in intent. First, we are very concerned with maintaining the spatial representation – i.e. foragers in a gridworld – rather than the discrete patch choices used by Harley, Regelmann, and many other modelers. Indeed, our results will show that the spatial nature of foraging is a critical aspect overlooked by many models. Furthermore, whereas Harley sought to show that IFD could be achieved with a simple rule, and Regelmann sought to show that deviations in competitive

abilities could affect the IFD, we are showing why undermatching should be expected more generally, and why Kennedy and Gray's (1993) meta-analysis found such a preponderance of undermatching in group foraging studies. Finally, although not discussed in detail in this article, our model also captures other emergent group-level phenomena such as population oscillations using the identical parameter values.

In our invisible model, each agent has memory for all cells in the 80 x 80 gridworld. At the beginning of a session, every cell in an agent's memory is initialized to a constant value. Thus, every cell looks equally attractive for action selection purposes. During a time step, each agent chooses a cell by estimating every cell's value and then using the softmax algorithm (Equation 2) to assign a movement probability to each cell. Whereas the visible model's agents decide between pellets by using current perceptual information, the invisible model's agents rely on their interaction histories with the environment. In particular, if an agent steps on a cell and receives a pellet, the cell's value receives a large boost in the agent's memory, and the neighboring cells are also boosted to a lesser extent. This neighborhood assumption leads to quicker learning, and its ecological validity arises from two phenomena: perceptual discrimination and generalization. On one hand, human foragers may not be making fine perceptual discriminations between cells, especially given the sheer number of cells. Thus, a forager may be rewarded at a location but not make a clear distinction between the rewarding cell and the rewarding cell's neighbors. Generalization, on the other hand, implies a more purposeful credit attribution to similar options. If one cell provides a reward, its neighbors may be likely to provide rewards too, because there is often underlying order in the environment. In any case, the boosted values in an agent's memory increase the probability of choosing those cells in the future. When an agent steps on a cell without receiving a pellet, however, a penalty is assessed to the cell and its neighbors. For both rewards and penalties, we set the relative strength of a neighbor's reward or penalty as an inverse function of its distance from the current cell. By exploring the environment and obtaining a history of rewards and penalties, each agent constructs a dynamic representation of the resource patches in the environment. Finally, much like the visible model, when an agent chooses a cell, a goal bias boosts the cell's value for the next action selection step,




Figure 1: Matching results for the visible and invisible models

thereby promoting probabilistic inertia towards a cell until it is reached. Equation 3 shows the additive calculation an agent makes for each cell, and the cell values are then input to the softmax algorithm to obtain selection probabilities.

Model Results

Undermatching

Figure 1 shows the matching results for the visible and invisible models at the three resource distribution levels of Goldstone and Ashpole (2004). The graphs show the proportion of agents in each pool (within a 25 unit radius of a pool's center) at a given time, after normalizing to exclude agents outside both pools. Each graph shows the averaged results from 30 trials, with 20 agents per trial. Importantly, although the parameter values are different for the two models, the same parameter values are used for all resource distribution levels of a model and for all subsequent manipulations discussed in this article. Our model results prior to normalization also match Goldstone and Ashpole's results prior to normalization. The matching results show that our agent-based models have captured the undermatching reported by Goldstone and Ashpole, and in fact our model results also show significantly more undermatching at the extreme 80/20 distribution levels. The same trials also exhibited the respective levels of population oscillations found by Goldstone and Ashpole, as well as high variance of agents' locations relative to resource locations and a comparable amount of switching between pools. We use these additional measures to confirm that the

models appear robust and reliable across several measures of group behavior, and we can now discuss additional simulations designed to discover the source of undermatching. Note that given Kennedy and Gray's (1993) meta-analysis and our earlier analysis of the Goldstone and Ashpole results, competitive abilities, travel costs, and the number of foragers could all be causes of undermatching. However, we can already eliminate competitive abilities as the explanation since the agents in our models share the same attributes.

Travel Costs

Travel costs between resource pools are sometimes used to explain undermatching (Baum & Kraft, 1998) due to the fact that foragers switch less frequently when the pools are farther apart, and therefore foragers obtain less information about the resources and make worse decisions. In order to test this explanation while trying to avoid possible confounds from changing the world size, we ran several simulations with agents starting inside the pools instead of random locations around the world. Figure 2 shows the matching results, averaged over 30 trials, from the visible model at the 80/20 distribution level, and results were similar for the invisible model and other distribution levels. Note that these experiments show less undermatching than Figure 1, but this is due to a more restrictive definition of being in a pool. For these simulations, we wanted to guard against possible matching biases caused by the pools' respective locations in the gridworld, so we placed the pools in diagonally opposite corners and restricted the pool definitions (20 units from the center) to avoid intersection




Figure 2: Agent starting locations

with the edges of the gridworld. In the left graph of Figure 2, we started all agents at random locations inside the 80% pool. In the middle graph, we started all agents at random locations inside the 20% pool, and in the right graph, we started 80% of agents inside the 80% resource pool and 20% of agents inside the 20% resource pool. Interestingly, all of these experiments result in the same undermatching asymptote despite differences in initial foraging behaviors. In addition, a more powerful observation can be made by examining the data prior to normalization. The 20% pool always has approximately 20% of agents, and undermatching arises because some agents leave the 80% pool and are outside both pools. The cost of switching does not seem to drive behavior. Of course, another argument against travel costs comes from the fact that Goldstone and Ashpole found undermatching in the visible conditions, despite participants constantly seeing the number of food pellets at each resource pool and thereby having relatively good information. Although we suspect that travel costs do indeed contribute to undermatching in many foraging situations, there appears to be another factor leading agents to selectively leave the plentiful pool and undermatch.

Number of Agents

We examined the effects of different numbers of foragers by running visible and invisible simulations at the 80/20 distribution level with 10, 20, and 30 agents, respectively. Once again, agents began in random locations, but here the food rate was held constant across conditions, with food dropped every 4/20 seconds (200 milliseconds). The constant food rate allowed us to gauge the effects due to the number of agents.

Figure 3 shows the matching results, averaged over 30 trials, for 10, 20, and 30 agents. These graphs are taken from the visible model, but the invisible model results are similar. Post-hoc ANOVA results indicate that 10 agents show significantly less ($p < .001$) undermatching than the 20 agents condition, and 20 agents showed significantly less ($p < .001$) undermatching than the 30 agents condition. Interestingly, none of the conditions are significantly different with respect to the 20% pool. Instead, the greater undermatching with respect to the large resource pool leads to more agents outside the pools. Thus, the number of agents -- or, equivalently, changes in food rate given a fixed




Figure 3: Number of Agents

number of agents -- does affect undermatching, but it is unclear why the effects take place.

Spatial Constraints

Our novel explanation for undermatching contends that the spatial environment of foraging critically determines how well foragers can distribute to the resources. Imagine a forager with a finite speed covering a finite territory. Now introduce five more foragers to that territory. As you add foragers, the reward rate for each forager correspondingly decreases. In the context of the Goldstone and Ashpole visible condition, a relatively small number of foragers can cover the two resource pools and easily pick up food soon after it is dropped every 4/N seconds, since the food rate is not high enough to tax their foraging abilities. The Gaussian food distribution enhances this effect, because only a few foragers are needed to pick up the majority of the food near the Gaussian centers, while a few additional foragers can dart to the food dropped on the peripheries. By this argument, most of the foragers are actually unnecessary, although they can certainly compete for resources. Thus, undermatching arises because only a limited number of foragers are required to pick up the available food in a given amount of time, and excess foragers simply increase indirect competition.

In the Goldstone and Ashpole experiments, both resource pools have equal variances, so the 20% pool can occupy as much space as the 80% pool. Although more food is dropped in the 80% pool, the spatial distribution means that there does not need to be an equal increase in the number of foragers to pick up that food. These considerations seem to be absent from the literature, and that may be due to the preponderance of models that treat foraging and probability matching as discrete choices between options, rather than often occurring in spatially instantiated situations. Furthermore, even empirical studies overlook this interplay of food rate, number of foragers, and spatial constraints. In Baum and Kraft's (1998) study of 30 pigeons distributing between two resource pools, the authors seem to erroneously conclude that food rate is unimportant. However, it appears that the authors did not use a sufficiently high feeding rate. Thus, a relatively small number of pigeons could eat the food regardless of the rate, because the rate never exceeded a few pigeons' capacities. As a result, pigeon distributions did not significantly change




Figure 4: Uniform Variance

because the rate increase did not correspondingly increase information.

Uniform Variances

Given the preceding argument for undermatching, we should expect significantly less undermatching if Goldstone and Ashpole had used uniform variance food distributions rather than Gaussian distributions, because the Gaussian distribution allows a lucky few to collect all the food at the center, greatly decreasing the distinguishing information between pools and requiring fewer foragers. Using our visible and invisible models, we conducted simulations with uniform variance food distributions. Figure 4 shows results from the visible model. The first number in each graph indicates the radius (for simplification, a centered square was used) of food droppings in the 80% pool, and the second number indicates the radius of food droppings for the 20% pool. Thus, the “16 vs. 8” condition creates an 80% pool that is four times larger than the 20% pool, so the probabilistic food distribution is reflected by pool size.

The graphs clearly support our undermatching explanation. When the pool variances are identical (8 vs. 8), the agents distribute themselves in approximately matching proportions. In these cases, food rate is the only factor that distinguishes between the two pools. In the 16 vs. 8 condition, the rarely observed phenomenon of overmatching occurs, and the explanation lies in the fact that the *densities* of the pools are equal, but the *coverage times* are unequal because the food rate is low. As each new piece of food is grabbed relatively quickly, foragers begin to converge as they chase new pellets. This convergence, along with the large pool size, can gradually increase the time it takes to reach a new pellet on a far side of the pool, and therefore the new pellet looks more attractive to migrators switching pools. Meanwhile, agents in the 20% pool have less area to cover and fewer pieces of food, so the pick-up time remains comparatively low. Finally, the reverse condition (8 vs. 16) leads to dramatic undermatching for similar reasons. The pick-up time in the 20% pool increases as agents converge, then new pellets dropped far away from the group attract additional foragers to switch from the 80% pool.

Conclusion

We have shown that agent-based models can produce complex group-level foraging behavior, and we have argued

for a novel explanation of undermatching: spatial constraints. We hesitate to use a single phrase for this phenomenon, because the effects truly arise from the interplay of spatial considerations with the food rate and the number of foragers. In our simulations, we showed that food distribution variances had a particularly large effect for this experimental set-up, but we suspect that other spatial issues – such as rugged terrain allowing different movement speeds – can lead to undermatching for similar reasons. In essence, we conclude that seemingly small changes in environmental parameters can have large effects on group foraging behavior due to the interaction effects of complex systems.

Acknowledgments

Support was provided by an NSF Graduate Research Fellowship to MER.

References

- Baum, W.M., & Kraft, J.R. (1998). Group choice: Competition, travel, and the ideal free distribution. *Journal of the Experimental Analysis of Behavior*, *69*, 227-245.
- Estes, W.K., & Straughan, J.H. (1954). Analysis of a verbal conditioning situation in terms of statistical learning theory. *Journal of Experimental Psychology*, *47*, 225.
- Fretwell, S.D., & Lucas, H.J. (1970). Ideal free distribution. *Acta Biotheoria*, *19*, 16-21.
- Gallistel, C.R. (1990). *The organization of learning*. Cambridge, MA: MIT Press.
- Gillis, D.M., & Kramer, D.L. (1987). Ideal interference distributions: Population density and patch use by zebrafish. *Animal Behavior*, *35*, 1875-1882.
- Godin, M.J., & Keenleyside, M.H.A. (1984). Foraging on patchily distributed prey by a cichlid fish (Teleostei, Cichlidae): A test of the ideal free distribution theory. *Animal Behavior*, *32*, 120-131.
- Goldstone, R.L., & Ashpole, B.C. (2004). Human foraging behavior in a virtual environment. *Psychonomic Bulletin & Review*, *11*, 508-514.
- Harley, C.B. (1981). Learning the evolutionarily stable strategy. *Journal of Theoretical Biology*, *89*, 611-633.
- Harper, D.G.C. (1982). Competitive foraging in mallards: Ideal free ducks. *Animal Behavior*, *30*, 575-584.
- Kennedy, M., & Gray, R.D. (1993). Can ecological theory predict the distribution of foraging animals? A critical analysis of experiments on the ideal free distribution. *Oikos*, *68*, 158-166.
- Magnuson, J.S. (1997). Does complex behavior require complex representations? *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society*, (pp. 472-477). Mahwah, NJ: Lawrence Erlbaum Associates.
- Regelmann, K. (1984). Competitive resource sharing: A simulation model. *Animal Behavior*, *32*, 226-232.