Conditions Enabling the Evolution of Inter-Agent Signaling in an Artificial World

James A. Reggia*,†
Reiner Schulz*
Gerald S. Wilkinson‡
Juan Uriagereka◇
Depts. of Computer Science*,
Biology‡, and Linguistics◇,
and Institute for Advanced
Computer Studies†
University of Maryland,
College Park, MD 20742 USA
{reggia,rschulz}@cs.umd.edu
gw10@umail.umd.edu
juan@wam.umd.edu

Abstract In the research described here we extend past computational investigations of animal signaling by studying an artificial world in which a population of initially noncommunicating agents evolves to communicate about food sources and predators. Signaling in this world can be either beneficial (e.g., warning of nearby predators) or costly (e.g., attracting predators or competing agents). Our goals were twofold: to examine systematically environmental conditions under which grounded signaling does or does not evolve, and to determine how variations in assumptions made about the evolutionary process influence the outcome. Among other things, we found that agents warning of nearby predators were a common occurrence whenever predators had a significant impact on survival and signaling could interfere with predator success. The setting most likely to lead to food signaling was found to be difficult-to-locate food sources that each have relatively large amounts of food. Deviations from the selection methods typically used in traditional genetic algorithms were also found to have a substantial impact on whether communication evolved. For example, constraining parent selection and child placement to physically neighboring areas facilitated evolution of signaling in general, whereas basing parent selection upon survival alone rather than survival plus fitness measured as success in food acquisition was more conducive to the emergence of predator alarm signals. We examine the mechanisms underlying these and other results, relate them to existing experimental data about animal signaling, and discuss their implications for artificial life research involving evolution of communication.

Keywords
evolution of communication, multi-agent systems, animal communication, animal signaling

1 Introduction

Some animals emit signals that provide information about their environment, such as indicating the presence of predators or nearby food, to others [17]. For example, several species of tamarins and marmosets as well as rhesus monkeys and chimpanzees emit calls upon discovering food and during its consumption [4, 7, 15, 18, 19, 49]. The use of alarm calls is also common. For example, ring-tailed lemurs give separate warning calls for aerial and terrestrial predators [35] and vervet monkeys use four phonically different alarm calls to indicate the identity of terrestrial, aerial, arboreal, or other predators [10, 38]. Understanding how resource recruitment and alarm signals might evolve requires identifying the conditions under which the benefit of signaling outweighs any cost. Advertising the location of food to potential competitors is unlikely to be profitable unless the food intake of the signaler is enhanced [16]. Similarly, alerting others to an approaching predator is unlikely to benefit the individual if it increases their risk of predation, as has been demonstrated for some species [40]. Close kinship provides the most plausible explanation for how costly animal signaling can evolve [16].

Motivated in part by the currently incomplete understanding of how animal communication about the environment can arise, there have been several past multi-agent models that have evolved or learned a shared communication system. In some cases the abstract signals/tokens used are ungrounded: they have no relation to an external environment [24, 29, 30, 41]. Others have simulated agents in environments with signaling of mating information [46], predators [1], food sources [44, 45], aggression [12], or robotic visual information [42]. In the research described here, which falls in the latter category of grounded signaling and is oriented toward vertebrate communication, we use an artificial environment in which a population of autonomous agents searches for food and tries to avoid predators. Our goal is to identify conditions in this simulated world under which simple signaling evolves among initially noncommunicating agents. Our research extends past modeling work on emergent signaling in two ways. First, we determine systematically how the density of agents and predators, and the amount and distribution of food, influence the evolution of signaling. These factors dictate the degree to which signaling is costly or beneficial to an individual. Most past models of emergent communication have not addressed this issue, although limited results from one recent study on food communication have indicated that these factors are of critical importance [44, 45]. Second, we determine how assumptions about simulated evolution, such as spatial constraints on mating and the extent to which fitness is based on food consumption, influence whether communication evolves. Past studies, including those cited above, have each used a single evolutionary process, but this has differed from study to study. Some models have been based on traditional genetic algorithm methods [22] that ignore spatial relationships between agents and use an explicit measure of fitness such as food/energy/resource consumption. In contrast, other models have based parental selection on physical proximity or solely on survival. We examine for the first time the impact of such choices on the ease with which communication can evolve. Of course, like past artificial life models of communication, ours is substantially simplified from reality to make computational investigation practical. For example, one important issue we do not examine here is the decoupling of sending and receiving of messages (see Discussion).

In the following, we first briefly summarize the artificial world and experimental methods used in this research. The results of varying the density of agents and predators, the distribution of food, and the selection of parents and placement of offspring are then described. Although some of these results appear intuitively plausible, others are less so and require substantial analysis to understand. Collectively they provide fairly concrete indications about the conditions under which the kinds of signaling considered here may arise. Further, they also indicate that the assumptions made about the simulated evolutionary process can substantially change the outcome, so that such assumptions must be carefully considered in future studies of emergent communication.

2 Methods

We first briefly describe the simulation environment created for this study and then consider the experimental methods used.

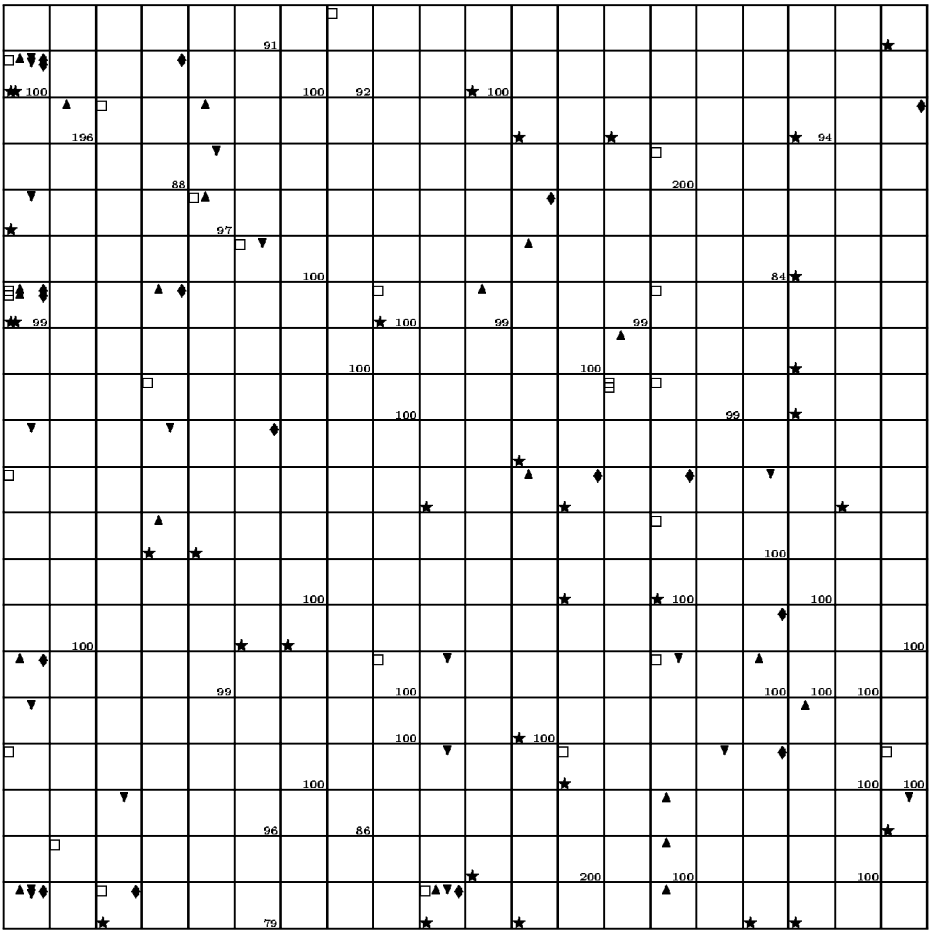

Figure 1. A 20 × 20 window into a 60 × 60 artificial world. Each filled triangle/diamond represents an agent from a specific class of communicating agents (same representation as in subsequent figures), each □ a noncommunicating (NC) agent, and each ⋆ a predator. If more than one predator or agent in the same class are present in a cell, they are indicated by overlapping symbols. The number at the bottom right corner of each cell is the number of food units there (blank if none). Thus, in the leftmost column, second cell down, there are five communicating agents of various types, one noncommunicating agent, two predators, and 100 food units. For illustrative purposes, the density of agents, predators, and food sites here is much higher than usually present in the simulations described below (800 agents, 300 predators, and 300 food sites in the entire 60 × 60 world).

2.1 Entities in the Simulation Environment

The simulation environment, or *world*, is similar to many environments used in previous artificial life research. Only a summary is provided here as a detailed description is available in [37]. The world consists of a bounded two-dimensional rectangular space of discrete cells, each cell representing a small region. At the start of a simulation, three types of entities are initially placed in random locations: agents, food sites, and predators (see Figure 1). Agents search for food sources and consume them when found. Agents flee from predators and can generally escape from a predator if it is detected in time. As agents can only "see" a small region in front of the cell they occupy in the direction they are oriented, predators may approach from other directions undetected. As long as the agent sees an approaching predator (it may not), it will escape unless it runs into another predator or a corner. Signals in this environment have the potential to benefit agents that "hear" them: Food signals can alert agents to food sources they might otherwise not find, and predator signals can warn agents of approaching but unseen predators so that escape is possible.

Food sites, initially placed in random locations, are static and passive ("plants"). Parameters govern the number of food sites and the initial amount of food per site. If the total amount of food present falls below a threshold due to its consumption by agents, one or more new randomly located food sites are automatically generated to restore food levels. The total amount of food present thus fluctuates due to its consumption and replenishment and may transiently exceed the total amount present initially.

Predators are simple, eternal, mobile, nonadaptive, nonevolving "machines" that hunt and kill agents. Predators exist in one of three states: quiescent, searching, and pursuing. Predators enter the quiescent state initially (to allow randomly placed nearby agents a chance to detect them and escape), following a kill, and following an unsuccessful chase. During the quiescent state the predator is completely idle and of no danger to agents. After a predetermined time, the quiescent state ends and predators automatically enter the searching state. Predators in the searching state move around in a quasi-random fashion, hunting for agents. Predators can see any agents within a certain predetermined distance in any direction. When a predator first sees one or more agents, it randomly selects one of the closest ones and enters the pursuit state with the selected agent as a target. In that state the predator repeatedly moves directly toward the selected agent until either it catches the agent (by landing on the cell occupied by the agent), it catches another agent inadvertently (by entering the cell of that other agent while pursuing a different one), or its pursuit time exceeds a predetermined maximum (the agent escapes). In all of these cases the predator subsequently enters the quiescent state. Predators can also hear signals issued by communicating agents within a set distance that is greater than the distance a predator can see. If a searching predator hears a signal, it changes its direction to move toward the source of the signal. In this sense, communication has a cost for agents in that it can attract predators.

Agents, the focus of this study, have more complex behaviors. Over time as they move around they construct a very limited "internal model" of their external environment based on what they can directly see and, if they communicate, also based on signals received from other nearby communicating agents. This internal model represents the existence and location of nearby predators and food but is limited by memory capacity and may contain inaccurate information (for example, the location of a food source that was consumed by other agents, or the location of a predator that has moved).

Agents in this world start where many past artificial life modeling efforts have ended: They have built-in, preprogrammed behaviors for avoiding predators and for seeking food. Such behaviors can lead to successful survival in the absence of communication. Our focus is on evolving communication that supplements this preprogrammed behavioral repertoire, leading to increased survival and fitness due to a better internal world model and cooperative actions. Accordingly, there are four classes or types of agents. These are designated NC for noncommunicating agents; F for agents that communicate about food only, emitting a *food signal* when arriving at a food site; P for agents that communicate about predators only, emitting a *predator signal* upon detecting a new predator; and FP for agents that communicate about both food and predators.

Each NC agent is represented as a table of information containing the agent's current location, direction of movement, direction of gaze, its current food stores, the total amount of food it has consumed since birth, memory of recent predators/food sites seen, and so forth. Food stores represent the current food reserves accumulated by an

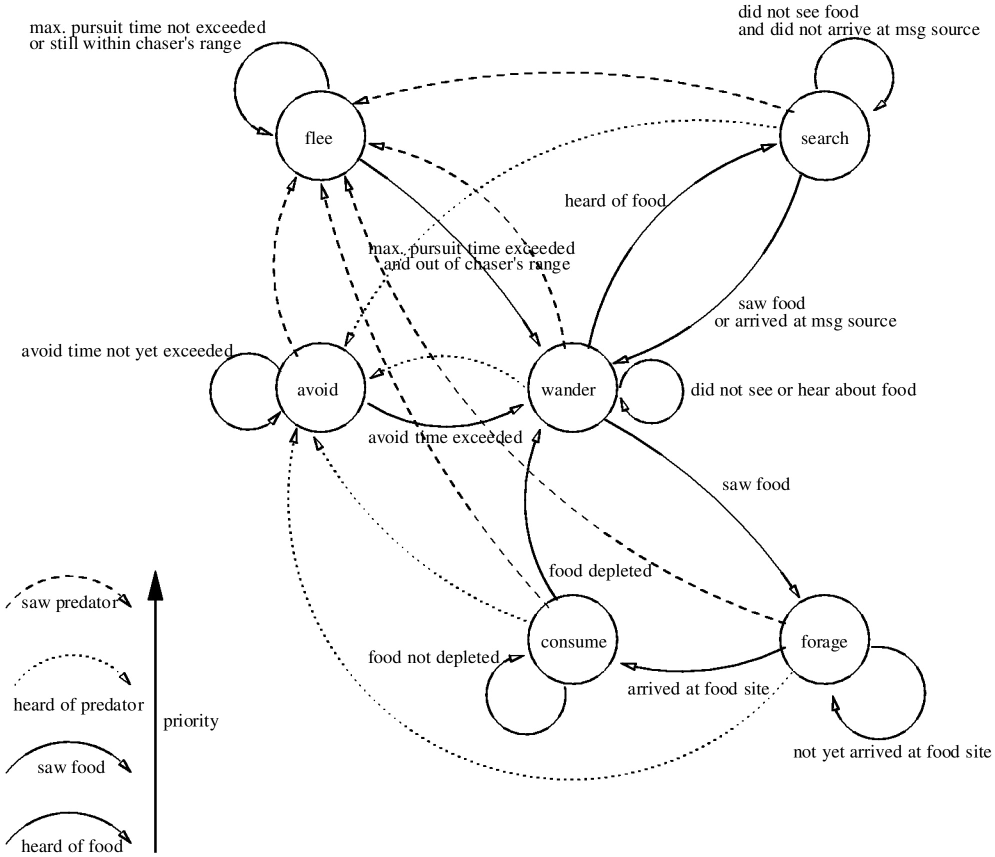

Figure 2. Automata model summarizing the behavioral states of agents. States are indicated by labeled circles and transitions by oriented arcs. Built-in transition priorities are depicted at the lower left. NC agents ignore signals.

agent. Each time an agent consumes a unit of food, its food stores are incremented. On the other hand, all agent food stores are decremented periodically, so an agent's food stores may rise and fall during a simulation. An agent will die if its food stores reach zero, if it is captured by a predator, or if it reaches a predetermined maximum age. Prior to starting a simulation, many parameters concerning agents can be adjusted to determine the agent population size, the distance an agent can see, the maximum food stores an agent can accumulate, how frequently food stores are decremented, the maximum age an agent can have, memory capacity, and so forth. As agents die during a simulation, enough new child agents are created (with other surviving agents serving as parent agents) at each time step so that the population size stays roughly constant.

Communicating agents incorporate all of the features of NC agents outlined above plus the ability to issue/receive signals. A single food or predator signal is issued when an agent first sees a nearby predator or arrives at a food site. The only other information associated with the signal is the location of the agent issuing the signal. A signal may be received only by other agents within a prespecified distance that is determined prior to starting a simulation. Agents may signal about food only, predators only, or both.

An agent's behavior is governed by which of six possible *states* it is in. These are described in more detail in [37] and summarized schematically in Figure 2. Each agent starts in, and may return to, the WANDER state, in which it has no current goal: It knows of no food sites or predators and therefore wanders aimlessly in the hope of discovering a food source. An agent will immediately enter the FLEE state if it sees a predator regardless of all other information or its current state (highest priority). In this state, the agent selects the closest predator that it sees, moves directly away from that predator, and then checks whether it is sufficiently far from the pursuing predator before terminating this state. The fleeing agent will escape unless it runs into a boundary ("cornered") or is within the range of another undetected predator that captures it. After flight, the agent returns to the WANDER state. The related AVOID state is entered if a predator-signaling agent (P or FP agent) detects a predator signal. The alerted agent moves directly away from the signal's source (the predator's location is uncertain but presumably close to the signal's source). The agent alternately looks backward to localize the predator if possible, and then forward to avoid running into another predator, as it moves.

The other three behavioral states relate to obtaining food. An agent will enter the FORAGE state if it is not fleeing or avoiding a predator, and if it has previously seen a food location. If so, it knows the exact location of the food site and moves directly toward it, looking in the direction of its movement. It is possible that previously seen food may no longer be present if it has already been consumed by other agents. If food is still present, the agent issues a food signal upon arriving at the food site indicating that it has found food (assuming the agent is a food-signaling F or FP agent) and enters the CONSUME state. In the latter state, if an agent's food stores are below their maximum capacity, the agent consumes one unit of food per iteration until full. It will remain at the food site until either all food is gone or it is forced to flee/avoid a predator. If the agent has filled its internal food store to capacity, it simply remains in the CONSUME state at the food site but does not consume any food until its food stores are decremented. When the food at that site is exhausted, the agent enters the WANDER state to determine its next action. Finally, if an F or FP agent has not seen food or predators, and is not avoiding a predator that it learned of from a signal, it will enter the state SEARCH to look for the closest food site about which it has previously received a signal. The agent in this state moves toward the original signal source (which was the food location) until a food source is seen or the agent arrives at the food location (assuming it does not see or hear of a predator in the interim). The agent then enters the WANDER state and other states (e.g., CONSUME) as appropriate.

2.2 Evolutionary Process

Each agent has a two-bit chromosome where the first bit indicates whether or not the agent sends and receives food signals, and the second whether or not it sends and receives predator signals. Thus, an agent with the chromosome 01 processes just predator signals, that is, it is a P agent. Chromosome values are assigned at the time an agent is born and, once assigned, are fixed. In the simulations described below, the initial population is always 100% randomly located NC agents. Of interest is the fraction of the agent population that evolves over time to communicate about food, predators, or both. During a simulation, whenever agents die and the actual number of agents present drops below the number present initially, the simulator automatically creates enough new agents (with new child agents coming from surviving agents at least 4 time steps old) to replace those that have died and places them in the world at the beginning of the next iteration. This allows total population size and density to be controlled, remaining essentially constant during a simulation without the need for substantial variations in model parameters.

New agents are evolved through a process of simulated natural selection and mutations (no crossover operation is done given the small chromosome). The evolutionary process is somewhat similar to what occurs with traditional genetic algorithms using tournament selection[1] but involves incremental replacement of dead agents rather than discrete generations and may take into account spatial relationships of parent/child agents. Whenever new agents are to be created, three steps occur: (a) selection of two parents, (b) generation of two children, and (c) placement of the children into the world. These steps are repeated as many times as necessary to restore the population to at least the target population size.

The first step in creating new agents is *parent selection*: Tournament selection is used to identify two parent agents to reproduce. In a tournament, a small set of candidate agents is picked, and the two of these candidate agents with the highest fitness are selected for reproduction (ties resolved arbitrarily). The fitness measure used is the current food stores possessed by an agent. This choice was motivated both by the consideration that animals must acquire food to produce offspring, with better-fed animals (larger, stronger, healthier, etc.) being expected to be more competitive for mating, and because food consumption has been used as a fitness measure in a number of past artificial life models. Two different tournament sizes, 2 and 10, are used in the simulations reported below to assess the impact of varying the amount and nature of competition during evolution. With a tournament size of 2, the fitness measure of an agent is irrelevant to selection: only an agent's "implicit fitness" in surviving starvation and predators matters. In contrast, with a tournament size of 10, after the 10 candidate parent agents are selected without regard for their fitness, the fitness measure is used to select the two most fit of the candidates to be the actual reproducing parents.

In addition to using two different tournament sizes, two different mutually exclusive mechanisms are used in selecting candidate parents for reproduction. With *random selection*, candidate agents are selected randomly and independently from the current population of agents. This initial selection of candidates is done without regard for agent fitness or the spatial location of agents, so the probability that an agent of a particular class (NC, P, etc.) is selected to participate in a tournament is the fraction of the population of that class at the time candidate selection is done. In contrast, with *spatially constrained selection*, only the first candidate agent is picked at random from the current agent population, while the remaining candidate agents are then the closest neighbor agents to the first one chosen (either 1 more or 9 more depending on tournament size; ties resolved arbitrarily). If the needed number of agents are not within a distance of 10 cells, fewer will be used. In the unusual situation where even a second agent cannot be found within this distance, this tournament is aborted and the original agent is rechosen. The spatially constrained selection of agents for reproduction, and the placement of their children near the parents (described below), would potentially tend to keep communicating agents closer together and thus could exert an influence on the evolution of communication.

The second step in the reproductive process is the actual *generation of children*. From the two parent agents selected as described above, two new child agents are created and added to the population. One new child agent has its genome initially set to that of one parent and the other new child agent has its genome initially set to that of the other parent. Rather than model recombination between food signaling and predator signaling loci, we use mutation to generate genotypic variation.[2] With probability p_{mf} the food bit in each child agent's chromosome is mutated (flipped; done independently for each child), and with probability p_{mp} (usually equal to p_{mf})

---

[1] The term "tournament selection" is used here in the sense that this term is usually used in the genetic algorithms and genetic programming literature [22, 26].

[2] With the small chromosome used here, an evolutionary process using mutation alone would be expected to produce results qualitatively similar to a process also incorporating recombination, so for simplicity no crossover operations were done in the simulations reported here. The broad issue of the relative value of recombination vs. mutation in the general area of evolutionary computation is substantial and controversial (see [27] for a discussion).

the predator bit in each child agent's chromosome is mutated (done independently for each child).

The third and final step in the reproductive process is the *placement of children*. The two new agents created by the above steps replace other agents that have died during the last iteration, and not the parent agents that remain in the world. If random selection of parents was used, then the two new child agents are always placed at random locations in the world, and we say that *random selection and placement* is used. If spatially constrained selection of parents was used, the child agent derived from each agent is always randomly placed within a radius of three cells of the parent agent from which it was derived, and we say that *spatially constrained selection and placement* is used. The initial direction of movement of a child agent is always random.

2.3 Experimental Methods

Unless explicitly noted otherwise, all simulations described below were done under the following conditions. The simulation environment was 60 × 60 cells in size, with edges being boundaries (agents could neither see nor move beyond edges). Predators and agents could see a distance of three cells, could hear agent signals a distance of six cells in any direction,[3] and generally moved a distance of one cell per iteration, horizontally or vertically or diagonally. Predators, and agents in the WANDER state, would move in the same arbitrary direction each time step with probability 0.9; on the other 10% of time steps they would randomly select a new direction for movements. Agents could see the contents of the cell they occupy and of cells in a 3 × 3 block of cells directly in front of them (vertically, horizontally, or diagonally, depending on direction of gaze). Agents had a maximum capacity for internal food stores of 30 food units. A newborn agent's initial internal food stores was 25 food units, and this was automatically decremented 1 food unit every 3 iterations. This allowed newborn agents substantial time to search for food site locations before their food stores became depleted. When at a food source location, an agent could consume one food unit per iteration, incrementing its internal stores a like amount, until it and any other agents present used up all of the food there. Agents had a maximum age of 1,000 iterations: It was generally unusual for an agent to survive so long in the simulations we examined, but if it did then it died from old age.

Simulations were done in which the number of agents, number of predators, number of food sites, and size of food sites were systematically varied. Within any one simulation, the values of these environmental parameters were held fixed. For each specific set of environmental parameter settings, four simulations were done using all possible combinations of random versus spatially constrained selection and placement and a tournament size of 2 versus 10. Each specific simulation was run 15 times, each time with a different random number generator seed (random numbers controlled or influenced the initial placement of agents, food sites and predators, the directions that agents or predators moved, the selection of parents during reproduction, whether or not mutations occurred, etc.). The results reported below for each simulation are the averaged values taken over these 15 runs, with variances indicated by error bars in graphs. Variances were usually quite small, giving confidence in the results, except for those parameter settings representing transitions between situations where simulations produced different qualitative results (e.g., for population sizes of 50, 75, and 100 in Figure 4B).

Simulations always started with 100% NC (noncommunicating) agents and were run for 100,000 iterations unless noted otherwise. A large amount of information about the

---

[3] Allowing agents to hear signals over a greater distance than they could see was intended to represent visual obscuring of food/predators by environmental conditions (e.g., foliage) with increased distance.

state of the world was collected automatically at periodic times during each simulation, the results being saved in files for subsequent analysis. The information measured includes the number of NC, F, P, and FP agents present, their average age and fitness, the number of births and deaths of each type of agent, how many deaths were due to starvation or predation for each type of agent, and so forth. Each of these measures was also averaged over the entire duration of simulations and over the last 30,000 iterations during which the fractions of each class of agent were more or less stable, typically fluctuating around some fixed value.

Our simulator is implemented in Allegro Common Lisp and compiled and run under UNIX. A single run of one simulation for 100,000 time steps with 200 agents takes approximately 30 minutes on a single dedicated workstation. To accommodate the large amount of computing time required by the roughly 11,000 simulator runs used to provide the results described below, typically 45 simulation runs were done simultaneously on 45 different workstations in a cluster of more than 60 SUN Ultrasparc 1 and 5 workstations, representing 15 runs each of three specific parameter settings.

3 Results

3.1 Baseline Simulations

To serve as controls, simulations were done in which no food and no predators were present, so there was no advantage to communication. Assuming a fixed mutation rate, the long-term expected fraction of the population in the baseline simulation that signals about food is 0.5, and the same would be expected for predator signaling. Since mutations of each gene occur independently, in any simulation it would be expected that each class of agent (NC, F, P, and FP) would ultimately form roughly 25% of the population.

Four simulations (60 runs of the simulator, since each simulation is run 15 times using different random numbers) were done in which the tournament size was either 2 or 10, and for each tournament size either random or spatially constrained selection and placement was used. In each simulation there were 200 agents (about 5.5 agents per 100 cells or 10 × 10 neighborhood), and a constant mutation rate of $p_{\mathrm{mf}} = p_{\mathrm{mp}} = 0.003$ was used. As expected, over the long term the fraction of each class of agent in the population fluctuated around 0.25. Figure 3A shows a representative example of this (tournament size of 2, spatially constrained selection and placement). Approximately 515,000 agents were created and evaluated during each simulation of 100,000 iterations with all agent deaths being due to starvation. This long-term “25% per class” distribution of agent types, where 50% of agents signal the detection of predators and 50% signal the discovery of food, thus represents the *baseline* or control values against which the results given below should be compared.

In contrast, in other conditions where food and/or predators are present along with a sufficient number of agents, signaling between agents provides them with a competitive advantage that leads to the evolution of communication between almost all agents. Figure 3B gives a typical example of the time course of such emergent communication and illustrates the two stages that usually appear over time during the evolution of signaling. In the first *transient stage*, rapid changes occur in the fraction of the agent population in each class. In Figure 3B, this occupies roughly the first 20,000 iterations, where noncommunicating agents give way to those that communicate about food only or predators only, and subsequently to agents communicating about both. In the second *steady-state stage*, each class of agents forms a fairly stable fraction of the total agent population. This stage begins at roughly iteration 20,000 in Figure 3B, with approximately 90% of the population being agents that communicate about both food and predators in this specific example.

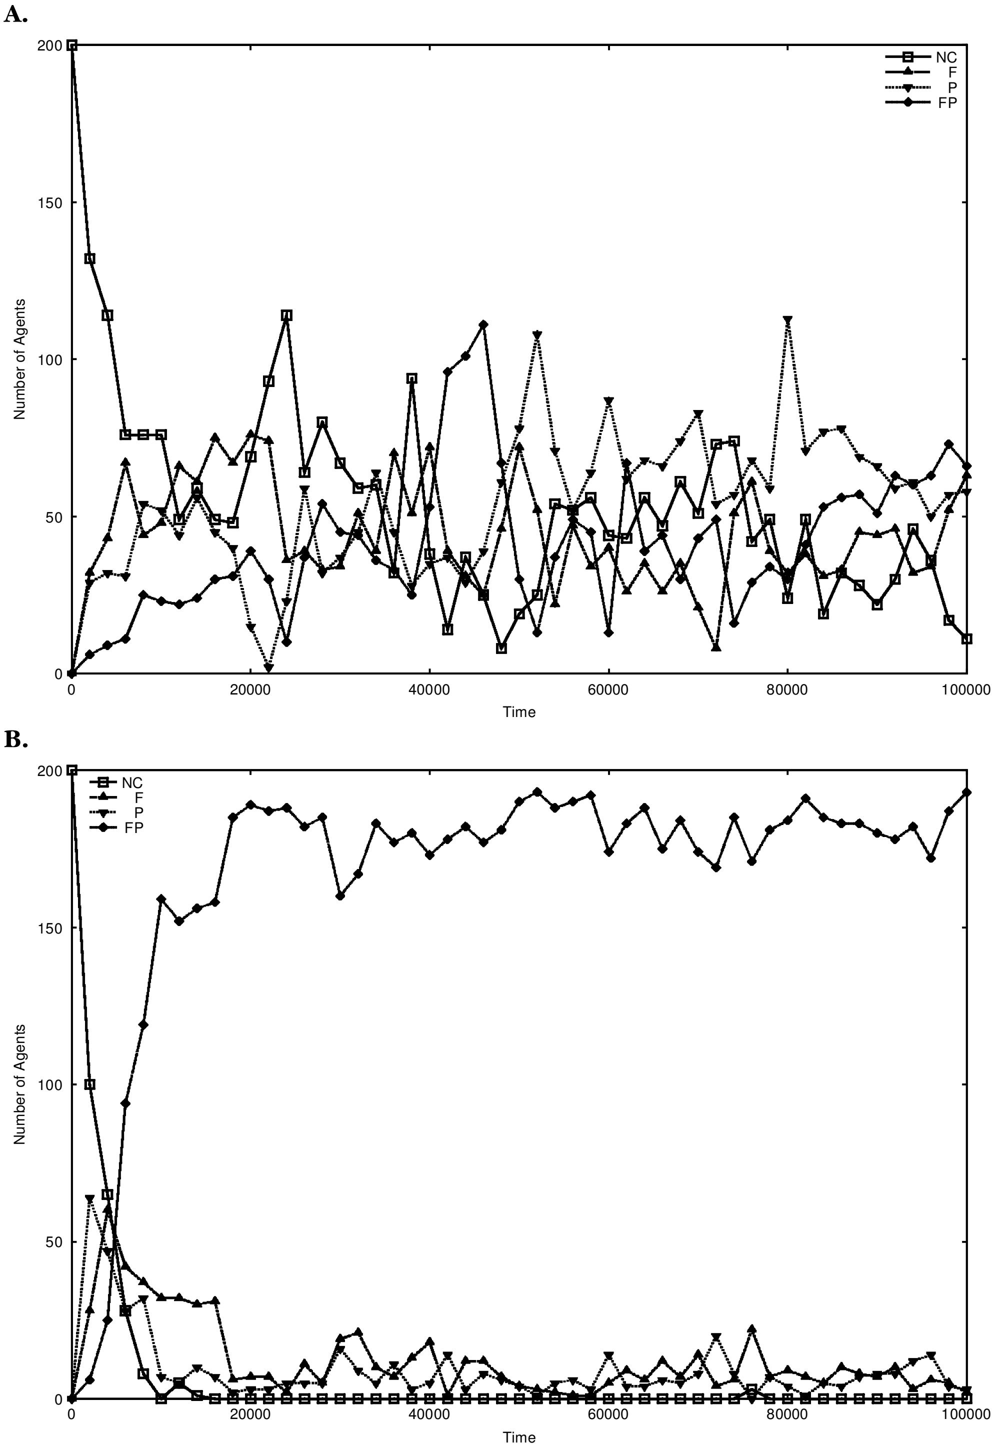

Figure 3. Numbers of the four classes of agents over time measured every 2,000 time steps (200 agents total, tournament size 2, $p_{mf} = p_{mp} = 0.003$). Each data point is the mean over 15 independent simulations having identical parameters. NC, noncommunicating agent; F, food-signaling agent; P, predator-signaling agent; FP, food- and predator-signaling agent. A) Baseline simulations with no food/predators (spatial selection and placement). B) Same parameters except 20 predators and eight food sites of 200 food units each present (random selection and placement).

3.2 Varying Agent Population Density

In the artificial world used here, agents can only hear one another over the limited distance of up to 6 cells. One would thus expect that below a certain density of agents, communicating agents would usually not detect each others' signals because they are generally quite far apart. It would be very difficult for communication to evolve in such a situation, regardless of how much it could potentially contribute to agent fitness, as communication would not be functionally effective. Analysis and measurement indicate that for a 60 × 60 world, with 25 randomly placed agents (less than 1 per 100 cells) present that communicate about the same topic (presence of predators or food), the expected distance to the *closest* other similar communicating agent is 6, whereas with 50 such agents the same expected minimum distance drops to 4. Thus one would anticipate that even in situations where communication was potentially highly advantageous, it would be unlikely to give a substantial advantage to 25 randomly located agents that can communicate up to a distance of only 6 cells. On the other hand, 50 such agents would generally be expected to have one or more other signaling agents within communication distance and thus should have a competitive evolutionary advantage.

As expected, simulations generally demonstrated that communication emerges more readily with higher population densities. This is illustrated in Figure 4A using random selection and placement in a world with 20 predators and 8 initial food sites at random locations, each with 200 food units. For a fixed world size of 60 × 60, as the total number of agents increased, the steady-state fraction of communicating agents tended to increase. When 50 total agents (1 per 70 cells) or fewer were present under these predator and food conditions, communication did not tend to emerge during the steady state beyond that of baseline control values. In contrast, for total population sizes of 100 or higher (roughly 3 agents or more per 100 cells), noncommunicating agents virtually disappeared, and roughly 90% of the agents communicated about both predators and food, the remaining 10% communicating about just food or predators. These results (Figure 4A) are generally consistent with the expectations outlined above. With a total of 50 noncommunicating agents initially present, about 25 agents (less than 1 agent per 100 cells) would be expected to evolve the ability to communicate about food, for example, due to chance mutations (baseline conditions). With only 25 such food-communicating agents initially present, they would in general be too separated to hear one anothers' signals as at best the closest neighboring agent would be expected to be at the outer fringe of possible communication distance, and the potential advantages of communication would not be realized. In contrast, with a total of 100 agents present, roughly 50 agents (about 1.5 agents per 100 cells) would be expected to communicate about food (baseline condition), and these would be sufficiently close together to hear and benefit often from one another's signals in general, so communicating agents would be expected to dominate, as was observed.

3.3 Varying the Selection Process

When spatially constrained selection of parents and placement of newborn child agents was used, results somewhat similar to those with random selection and placement were obtained. However, substantially more communicating agents could emerge at the lowest agent densities, as illustrated in Figure 4B. Here the simulations are identical to those in Figure 4A except that spatially constrained rather than random selection and placement occurred. For example, with only 25 agents present (leftmost in Figure 4B), about 75% of agents ultimately signaled predator detection, and 70% the discovery of food, while the number of NC agents at steady state dropped to less than 5%. This was because the spatially constrained selection and placement process led to local increases in density of communicating agents, and because the signaling of food

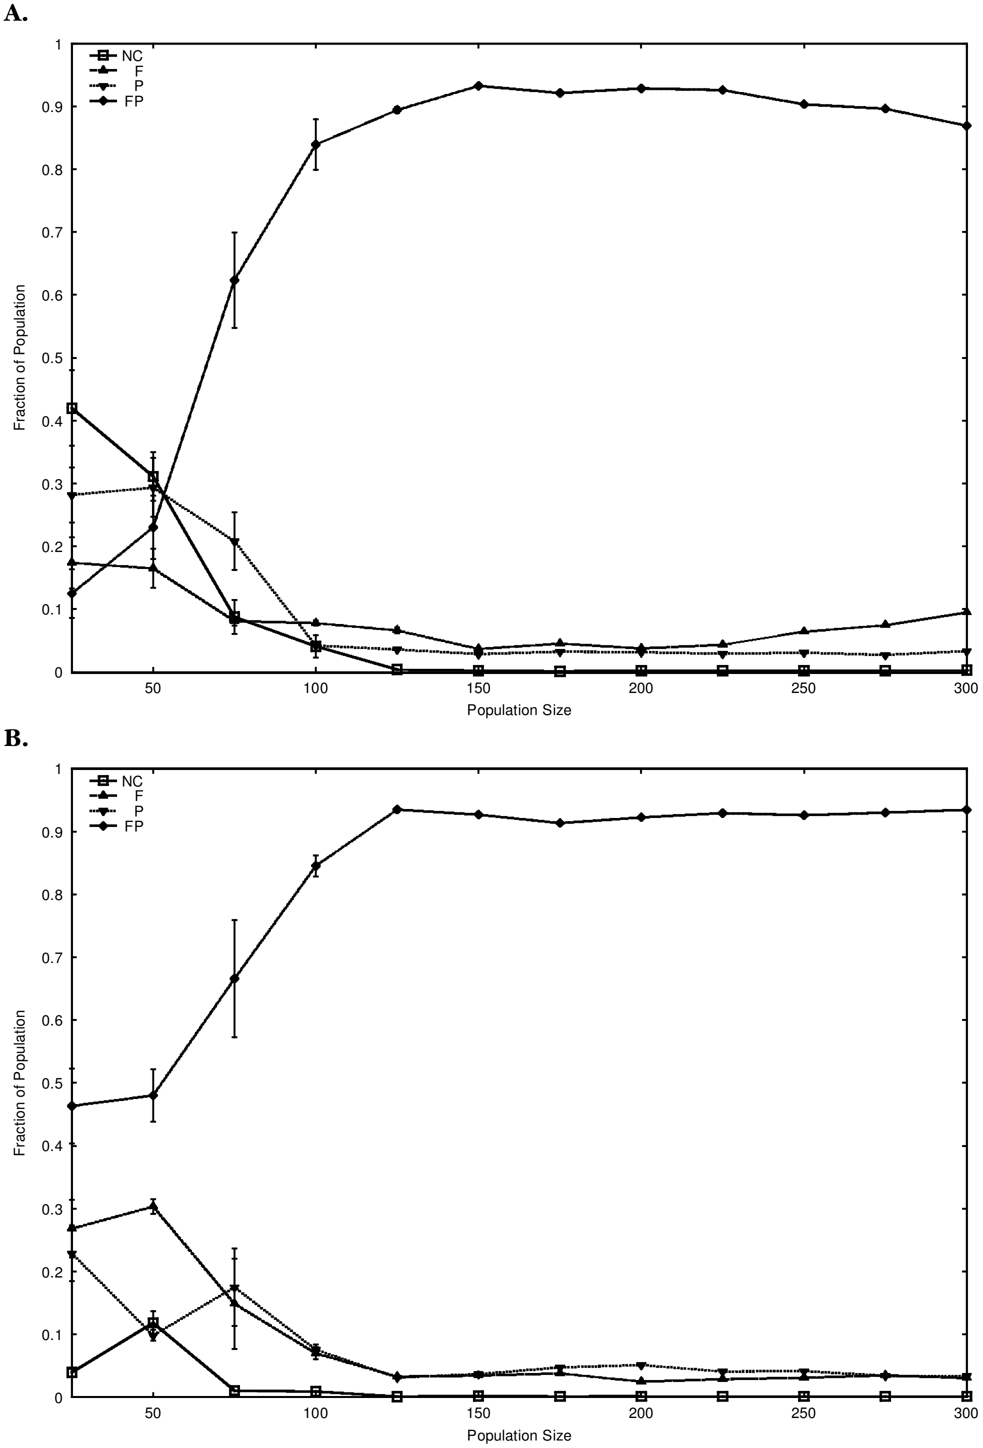

Figure 4. Fraction of agent types averaged over last 30,000 time steps of simulations versus size of population. Each data point is the average of results from 15 independent simulations having identical parameter settings (error bars are variance). Tournament size 2, mutation probabilities 0.003, 20 predators, eight food sites of 200 food units each. A) Random selection and placement. B) Spatially constrained selection and placement.

Table 1. Influence of tournament size on predator signaling

Tournament Size	Fraction of population		Mean age		Mean fitness		Starvation rate[a]		Predation rate[a]		Fraction fleeing	
	NC	P	NC	P	NC	P	NC	P	NC	P	NC	P
2	.05	.95	38	47	12	14	.0047	.0065	.0115	.0090	.14	.33
3	.78	.22	49	45	15	13	.0053	.0061	.0097	.0090	.16	.28
5	.92	.08	50	40	15	12	.0053	.0057	.0096	.0092	.16	.23
10	.96	.04	50	36	15	10	.0054	.0058	.0095	.0091	.16	.21

[a] Rate = (total number agents starving/killed during 30,000 iterations) ÷ [(mean number agents present) × 30,000]

caused communicating agents to cluster around food sites more than NC agents did. As a result, communication could become more effective at a lower global density of agents. Further, this clustering of food-signaling agents meant that during reproduction, the probability that the second, nearby agent to a signaling agent was also a signaling agent was increased due to spatially constrained selection, as discussed later.

Another factor influencing the evolutionary process is the tournament size. All of the results above are for a tournament size of 2. With this tournament size, the fitness function based on agent food stores plays no part in selection for reproduction: All agents that avoid death by starvation or predation are equally likely to be selected as a parent, so all that counts in an evolutionary sense is survival and not food stores. With larger tournaments of size 10, the fitness function comes into play in a major way, as from the 10 candidate parents initially selected in a tournament, only the 2 most fit in terms of food stores are used to create offspring. Candidate parent agents are thus strongly favored if they are successful at food acquisition.

The influence of the fitness function with larger tournament sizes had a dramatic effect on the evolution of widespread communication. Repeating the same simulations as in Figure 4 but now with a tournament size of 10 results in emergence of agents that communicate only about food and not predators, when the total agent population is sufficiently dense. As illustrated in Figure 5, with a total agent population size of 250, or roughly 7 agents per 100 cells, over 90% of the population ultimately does food signaling but less than 10% signals predator discovery, in spite of the presence of 20 predators. This means that not only was food communication rewarded, but also predator communication was effectively *punished* (recall from the baseline studies, 50% of agents would be expected to communicate about predators even if predator communication was a neutral factor in evolutionary terms). Spatial selection and placement did not facilitate predator communication in this case.

A series of simulations was undertaken with different tournament sizes (2, 3, 5, 10) to examine further when and why increased tournament size suppressed communication about predators. These simulations all had the same settings for other parameters (200 agents, all initially noncommunicating agents, 20 predators, 8 food sites of 200 food units each, random selection and placement). No food communication was permitted, so the agent populations either evolved to communicate about predators only or not to communicate at all. The results are given in Table 1, where each row represents the averaged result of 15 simulations having the same tournament size. Each pair of columns after the first column gives various steady-state measures averaged over the last 30,000 iterations of simulations, for NC and P agents only.

Table 1 shows that some suppression of signaling about predators occurs for *any* tournament size above 2, that is, for any simulations in which the current food stores (explicit fitness measure) come into play in selecting parents for reproduction. The

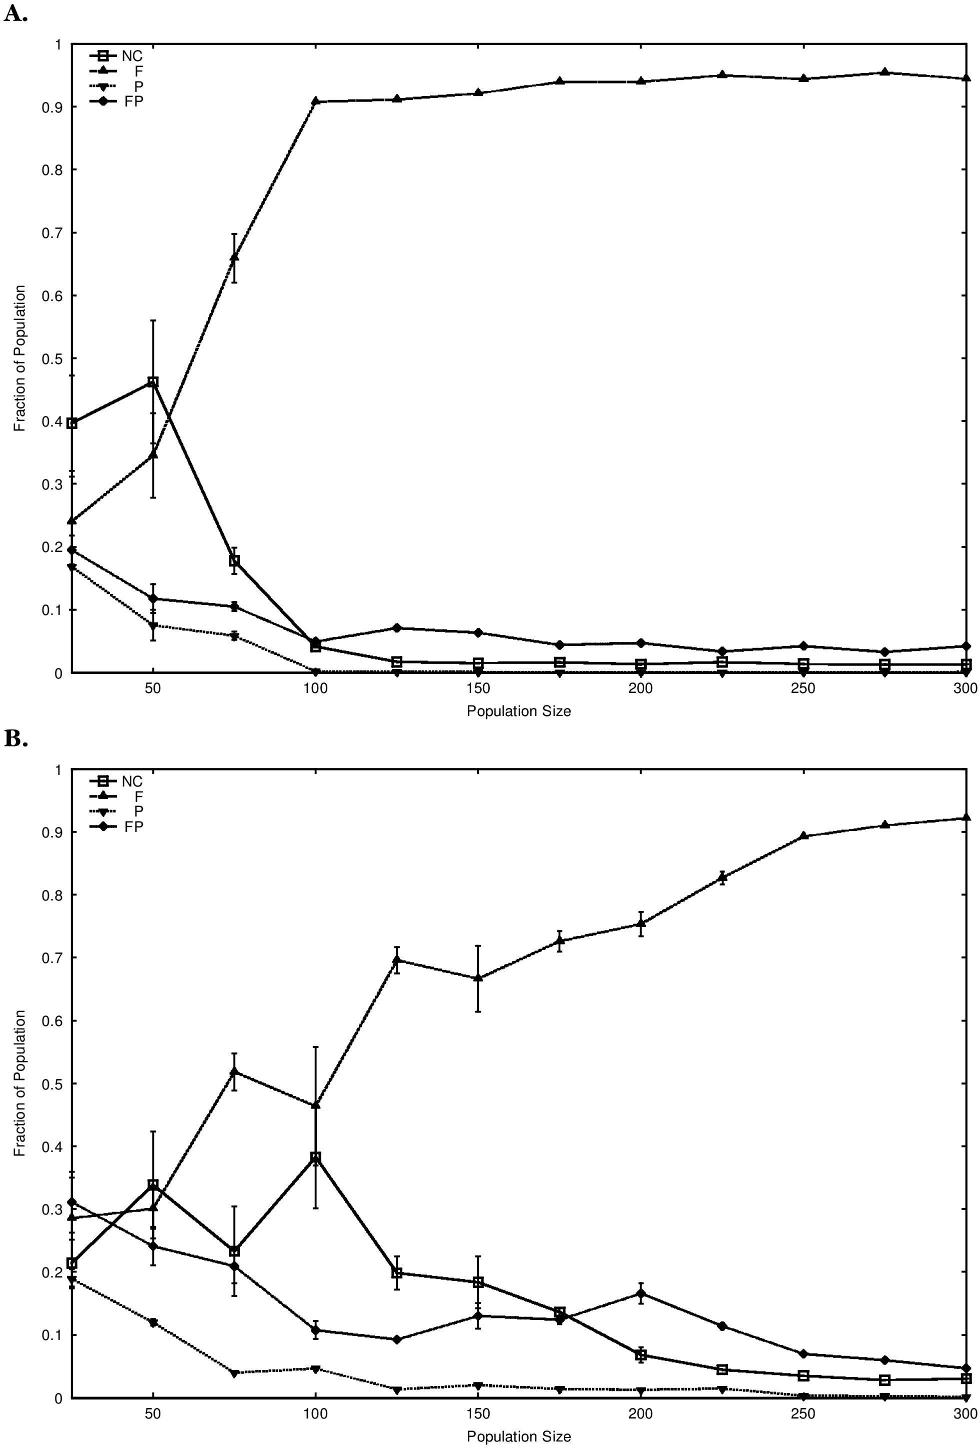

Figure 5. Identical simulations as in Figure 4 except that a tournament size of 10 rather than 2 is used. A) Random selection and placement. B) Spatially constrained selection and placement.

results in Table 1 provide insight into why predator signaling is penalized like this for larger tournament sizes. In general, P agents have a lower death rate due to predation than NC agents. While this comes with a cost—a higher starvation rate due to more time fleeing from predators and thus less time spent consuming food—the increased survival due to signaling about and thus avoiding predators is more potent. Since fitness in terms of food stores does not enter into the selection process with a tournament size of 2, the increased survival times of P agents cause them eventually to dominate the population. In contrast, for tournament size 10 and other values above 2, although predator signaling still helps reduce predation deaths among P agents, this is again at the expense of food consumption, so fitness of P agents tends to be less than that of NC agents (see Table 1). The fitness measure, favoring the NC agents, provides a strong bias built into the parent selection process, which is the dominant factor with larger tournament sizes, making P agents much less likely to be selected for reproduction that would produce P child agents. The more explicitly fit NC agents thus persist over the long term.

3.4 Variable Number of Predators

A reasonable hypothesis is that with larger numbers of predators present, more selection pressure will exist and encourage the evolution of agents that give warnings of nearby predators. As seen in Figure 6A, this trend was certainly the case with smaller numbers of predators. With random selection and a tournament size of 2, when no predators are present (leftmost on graph), roughly 50% of agents are FP and 50% F, indicating that communication about predators is neither rewarded nor punished in an evolutionary sense. When predators are present, as the constant number of predators present in simulations increases from 0 to 12, the latter representing about 1 predator per 300 cells, there is a substantial increase in the fraction of the population ultimately communicating about predators up to 90% or more. Further increases in number of predators beyond 12 did not make a significant difference. Essentially the same results were obtained using spatial selection and placement (not shown for brevity).

Substantially different results were obtained when the same experiments were repeated with a tournament size of 10 so that fitness based on food stores is an important factor. In this case, as shown in Figure 6B, the fraction of agents evolving to communicate about predators dropped to 10% or less, regardless of how many predators were present. As explained above, this reflects the cost of excessive flight times of agents, diminishing their food stores and thus reproductive fitness, in the context of larger tournament size. When spatial selection and placement was used with a tournament size of 10, results similar to those in Figure 6B were obtained (not shown for brevity).

3.5 Variable Amounts of Food

Determining conditions under which food signaling will evolve among initially noncommunicating agents in the absence of predators is more complex, depending not only on the total amount of food present but also its distribution. For example, for a fixed total amount of food (1,600 units here), Figure 7A shows that as the number of randomly placed food sites is increased, three types of population effects can be observed in the absence of predators. First, with just one very large food site, there is little or no substantial effect on the evolution of communication (roughly 50% NC and 50% F agents ultimately coexist). Second, with roughly 4–16 total food sites, which is about 1–4 food sites per 1,000 cells, that are still relatively large (> 100 food units each), communication about food arises early and communicating F agents essentially completely replace noncommunicating NC ones. Under these conditions F agents are better able to find food, and there is plenty of food for multiple agents to share at each site. Third, with a greater number of food sites (roughly 40 sites or more, corresponding

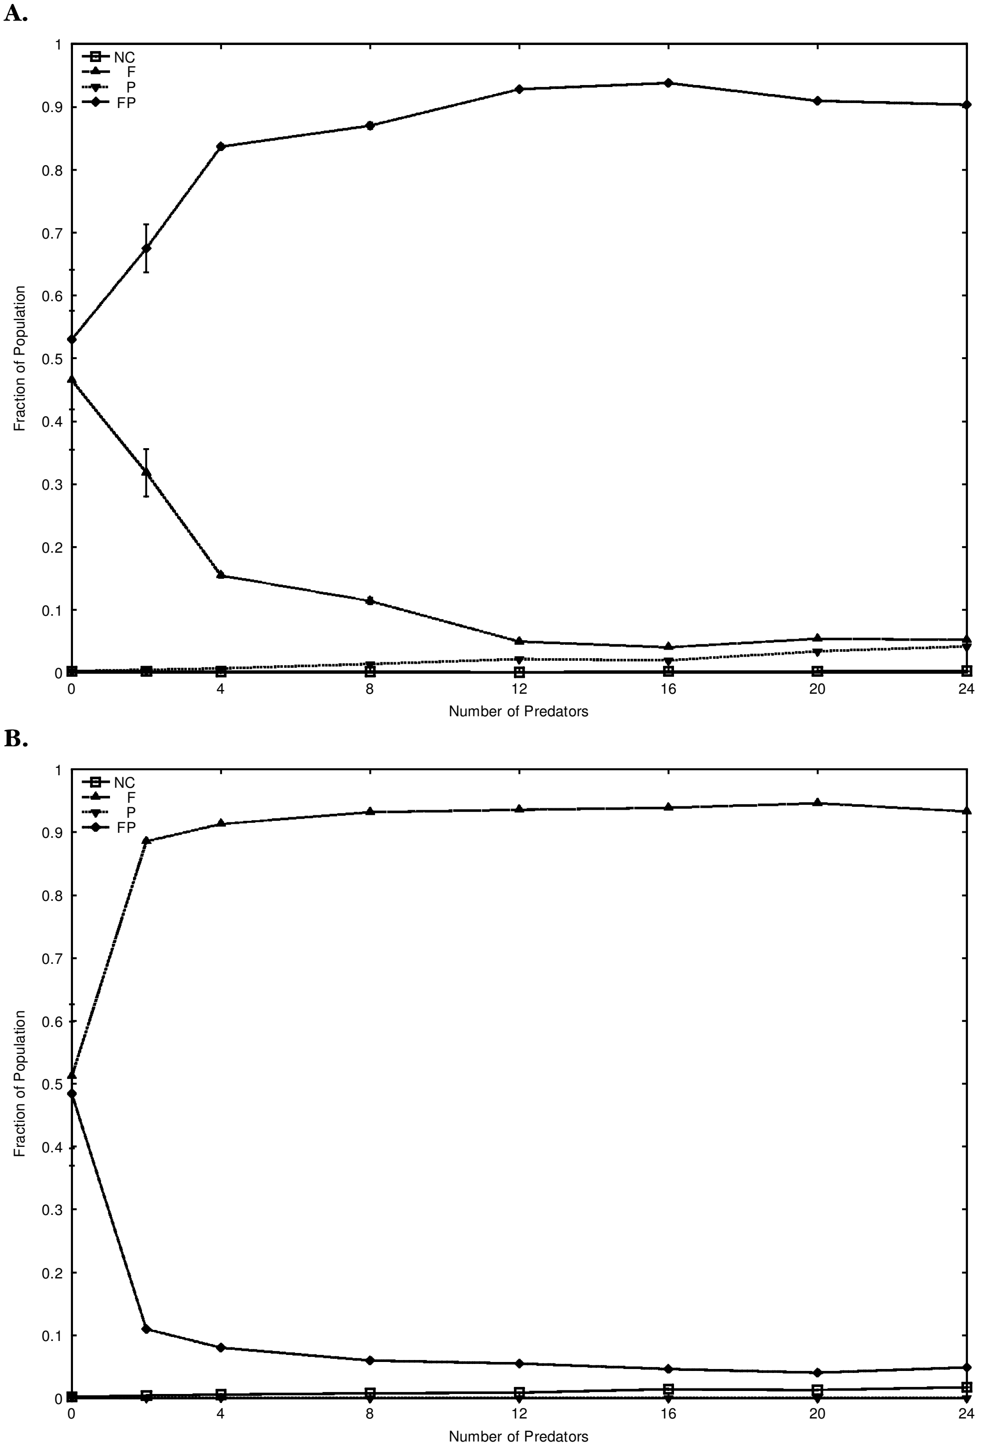

Figure 6. Fraction of agent population with each genotype averaged over last 30,000 time steps of simulations with different numbers of predators present. Parameter settings: random selection and placement; eight food sources of 200 food units each; mutation probabilities 0.003. A) Tournament size 2. B) Tournament size 10.

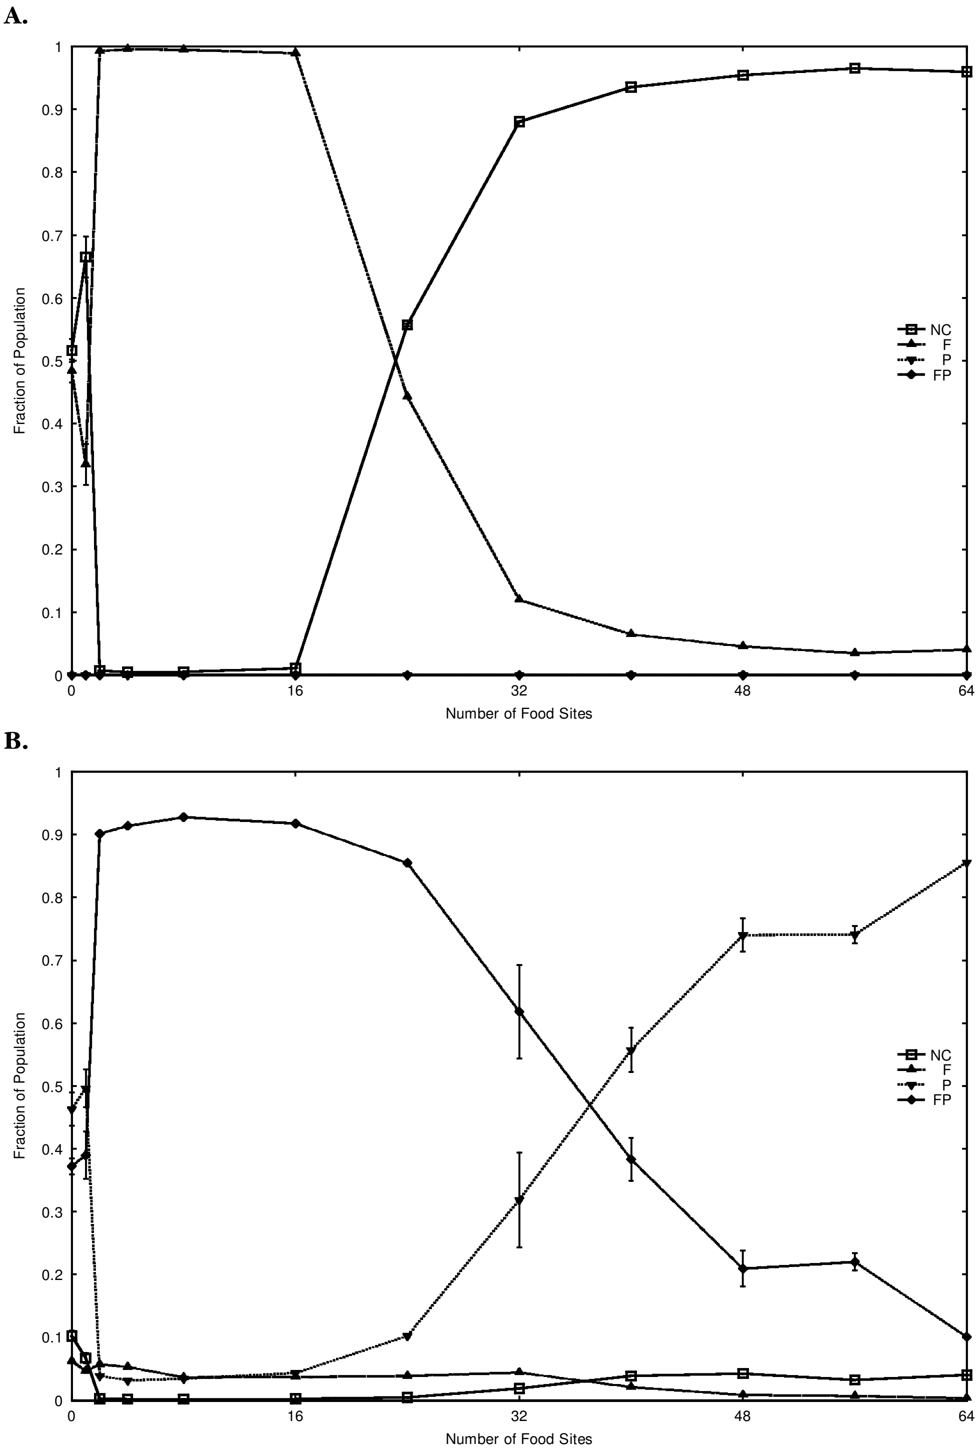

Figure 7. Fraction of agent population with each genotype averaged over last 30,000 time steps of simulations with different distributions of food present. Total food present is held constant at 1,600 units. Setting: random selection and placement; tournament size 2. A) No predators present: $p_{\mathrm{mf}} = 0.003$; $p_{\mathrm{mp}} = 0.0$. B) 20 predators present: $p_{\mathrm{mf}} = p_{\mathrm{mp}} = 0.003$.

to 1 food site or more per 100 cells), communicating about food not only is no longer advantageous but has even become disadvantageous, so that food communication is suppressed. In the latter case, the amount of food per site is quite limited. For example, with 50 food sites, at the time a food site is created it has 32 food units, and this rapidly decreases as soon as it is discovered by one or more agents. In this case, signaling the presence of food often attracts multiple other agents who compete in consuming the limited amount of food that is present, whereas noncommunicating NC agents will be more dispersed, likely to discover food on their own since there are so many sites, and much less likely to have to share any food they discover. Although the results given in Figure 7A are for random selection and placement and a tournament size of 2, almost identical results are obtained using spatial selection and placement, and in both cases also when a tournament size of 10 is used.

Further insight into how the distribution of food influences the evolution of food signaling can be obtained by keeping the number of food sites constant while varying the amount of food per site in the absence of predators. Noting the three ranges of number of food sites in Figure 7A that were neutral, advantageous, and disadvantageous for the emergence of food signaling, the effects of varying the amount of food per site were examined when a fixed number of 1, 8, or 56 food sites was present. For each fixed number of food sites present, the ultimate prevalence of food signaling as the amount of food per site was varied was examined under the usual four conditions: when selection and placement was random or spatially constrained, and in each of those cases when a tournament size of 2 or 10 was used.

When only a single food site is present, then regardless of the amount of food that is at the site (varied systematically from 0 to 1,600 units in different simulations, but held fixed in any given simulation), the conditions were neutral with respect to evolution of food communication. In general, 50% ± 5% NC agents and 50% ± 5% F agents were present at the end of simulations. This was true both when agent selection and placement was random and when it was spatially constrained, and regardless of whether a tournament size of 2 or 10 was used. The class of agents favored slightly, NC or F, varied arbitrarily, and neither class was favored overall. Clearly, a single isolated food site exerts minimal influence on the evolution of food signaling in this setting.

When eight food sites are present, so there is about one food site per every 450 cells, a very different and more complex pattern emerges as the larger number of food sites is adequate to exert an impact on the evolution of communication. Multiple, at times conflicting factors now operate to influence survival and parent selection. Which factors will dominate is quite difficult to predict a priori. When a relatively small amount of food per site was present, often food signaling clearly *impaired* the survival of agents, so food signaling agents only ever became a minority of the population. Figure 8A gives an example of this when random selection and placement and a tournament size of 2 are used: Food signaling agents only became about 25% of the population when roughly 5–20 food units per site are present. With random selection and placement but a tournament size increased to 10, as seen in Figure 8B, this effect became even more pronounced: as little as 10% of the population could become food signalers, and this effect occurred over a broader range of amounts of food per site (roughly 10–50 food units per site). In summary, as illustrated in the first two rows of Table 2, when random selection was used with eight small food sites, F agents communicating about food tended to be suppressed: They did not reach the expected 50% baseline fraction of the population. This suppression occurred because food signaling had a cost even in the absence of predators. An agent that signaled food discovery attracted other agents to that site, so the signaling agents competed for and recovered less food from the limited amount of food present, particularly when several communicating agents were attracted to the same site (their subsequent signaling perhaps attracting still others).

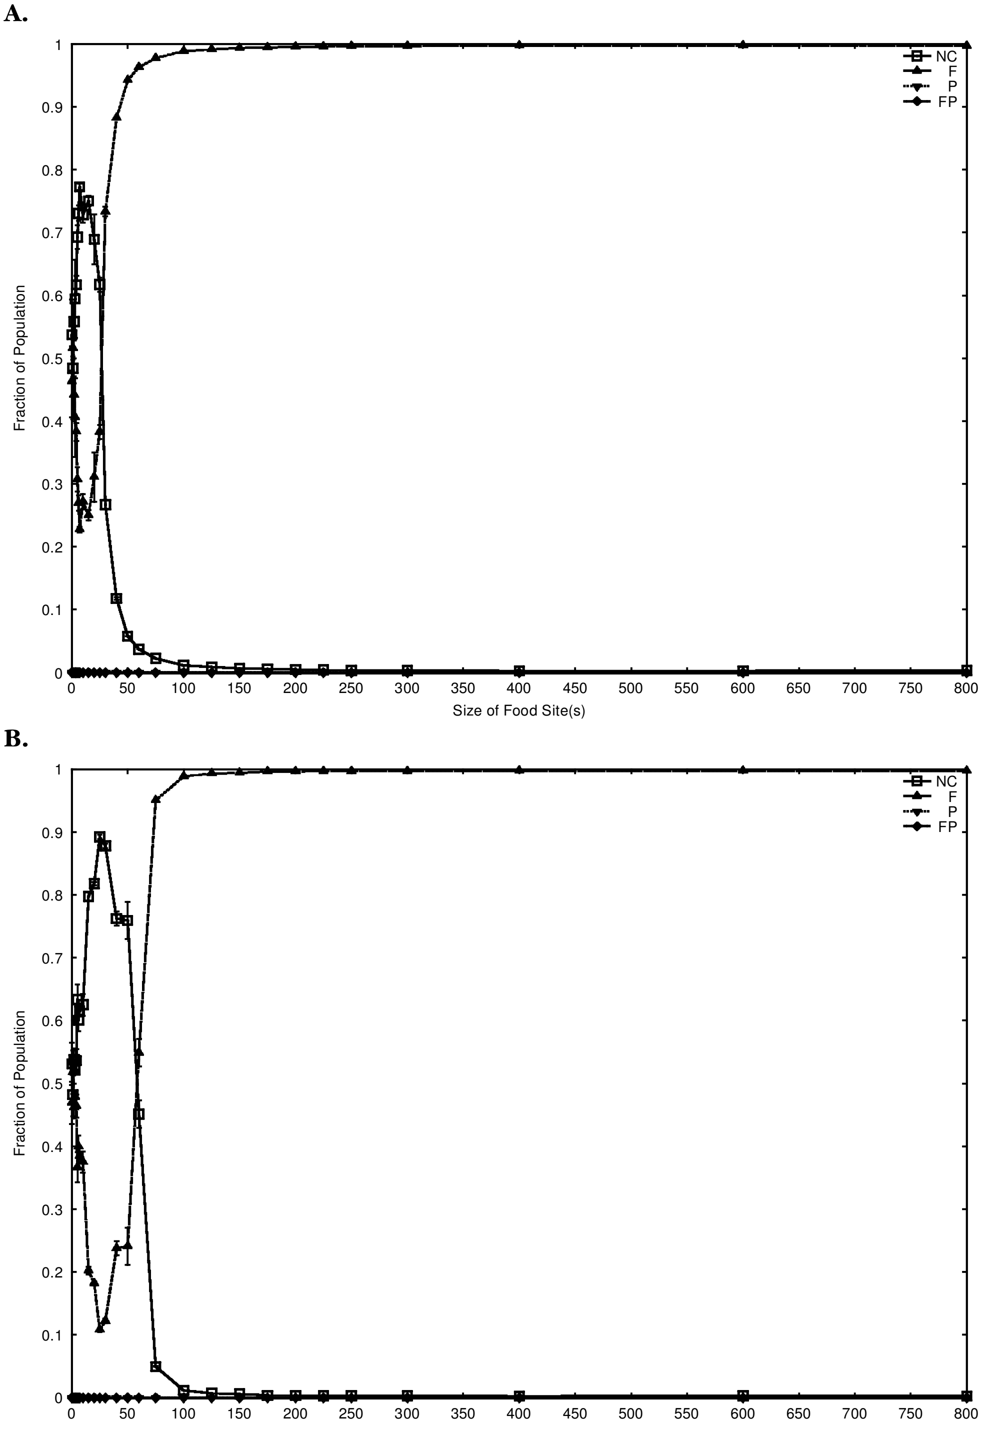

Figure 8. Fraction of agent population with each genotype averaged over last 30,000 time steps for simulations with different size food sites. Eight food sites are present in all simulations. Random selection and placement is used and no predators are present ($p_{\mathrm{mf}} = 0.003$; $p_{\mathrm{mp}} = 0.0$). A) Tournament size 2. B) Tournament size 10.

Table 2. Measures with eight food sites, 10 units of food per site (averaged during last 30,000 iterations of 15 runs)

Selection/ placement	Tournament size	Fraction of population		Mean age		Mean fitness		Nearby agents[a]	
		NC	F	NC	F	NC	F	NC	F
Random	2	.68	.32	43.8	42.4	11.9	11.7	7.4	5.8
Random	10	.76	.24	42.5	33.3	11.4	9.0	8.4	6.9
Spatial	2	.11	.89	38.1	42.3	10.1	11.8	5.0	11.7
Spatial	10	.34	.66	38.3	39.7	10.0	11.0	7.8	11.4

[a]Mean total no. of other agents within radius 6 of an agent of indicated type, per agent

Thus, under these conditions agents signaling about food actually obtained less food than noncommunicating agents, starved more often, and died younger, on average (see Table 2).

In contrast, with eight food sites where spatially constrained rather than random selection and placement was used, food-communicating agents could outnumber noncommunicating agents even when the amount of food per site was small. This is illustrated in Figure 9A for spatial selection and placement and a tournament size of 2. Here food-signaling agents became roughly 75% of the agent population when there were 5–20 units of food per site. In the case with 10 units of food per site, F agents formed almost 90% of the population, living longer and having a higher fitness than noncommunicating agents (row 3 of Table 2). This is quite a difference relative to random selection and placement with the same tournament size of 2 (compare with row 1 of Table 2). The difference in population sizes arose primarily because, as noted earlier, the food-signaling agents are more spatially clustered together in simulations than noncommunicating agents, as food signaling attracts communicating agents to one another at food sites (see nearby agents in Table 2).[4] Because of this, with spatially constrained but not random selection and a tournament size of 2, whenever the first parent selected for reproduction is a food-signaling agent, the likelihood that the *second* parent selected (the closest other agent) is also a food signaling agent rises substantially.[5] Thus, birth of food-signaling agents occurs at a faster rate than would be expected by chance (random selection) based on fraction of the overall population. This increase in births of communicating F agents with spatially constrained selection more than compensated for the disadvantage of food-signaling agents in the presence of a few small food sites. However, it is critically dependent upon tournament size. As can be seen in Figure 9B, with a tournament size increased to 10, only for the very smallest food sites (10 or fewer units of food per site) did food communicators dominate, with F agents becoming about 75% of the population. With a tournament size of 10, more noncommunicating agents were drawn into the process of selecting the second parent agent when the first was an F agent, and their sometimes higher fitness led to their more frequent selection for reproduction, so the deleterious effects of signaling the presence of small food sites could still predominate.

---

[4] This can be shown in various ways. For example, if one simulates a pure population of 200 noncommunicating agents without mutations, always replacing dying agents with new NC agents, the mean number of agents within the spatial selection range is 9.0 with random selection and placement, and 9.4 with spatially constrained selection and placement. With a pure population of 200 F agents, the same values are 11.1 and 12.3, respectively.

[5] This was true regardless of the fraction of the population that is F agents. For example, if rather than evolving agents one replaces dying agents so that the population is held fixed at 75% NC agents and 25% F agents, then it is found that 31% of closest neighbors to F agents are also F agents, a value higher than the expected fraction of 25%, due to F agents being attracted together by signaling. Since fitness did not enter into parent selection with a tournament size of 2, the second parent of an F agent was more often an F agent than would be expected based on their fraction of the overall population.

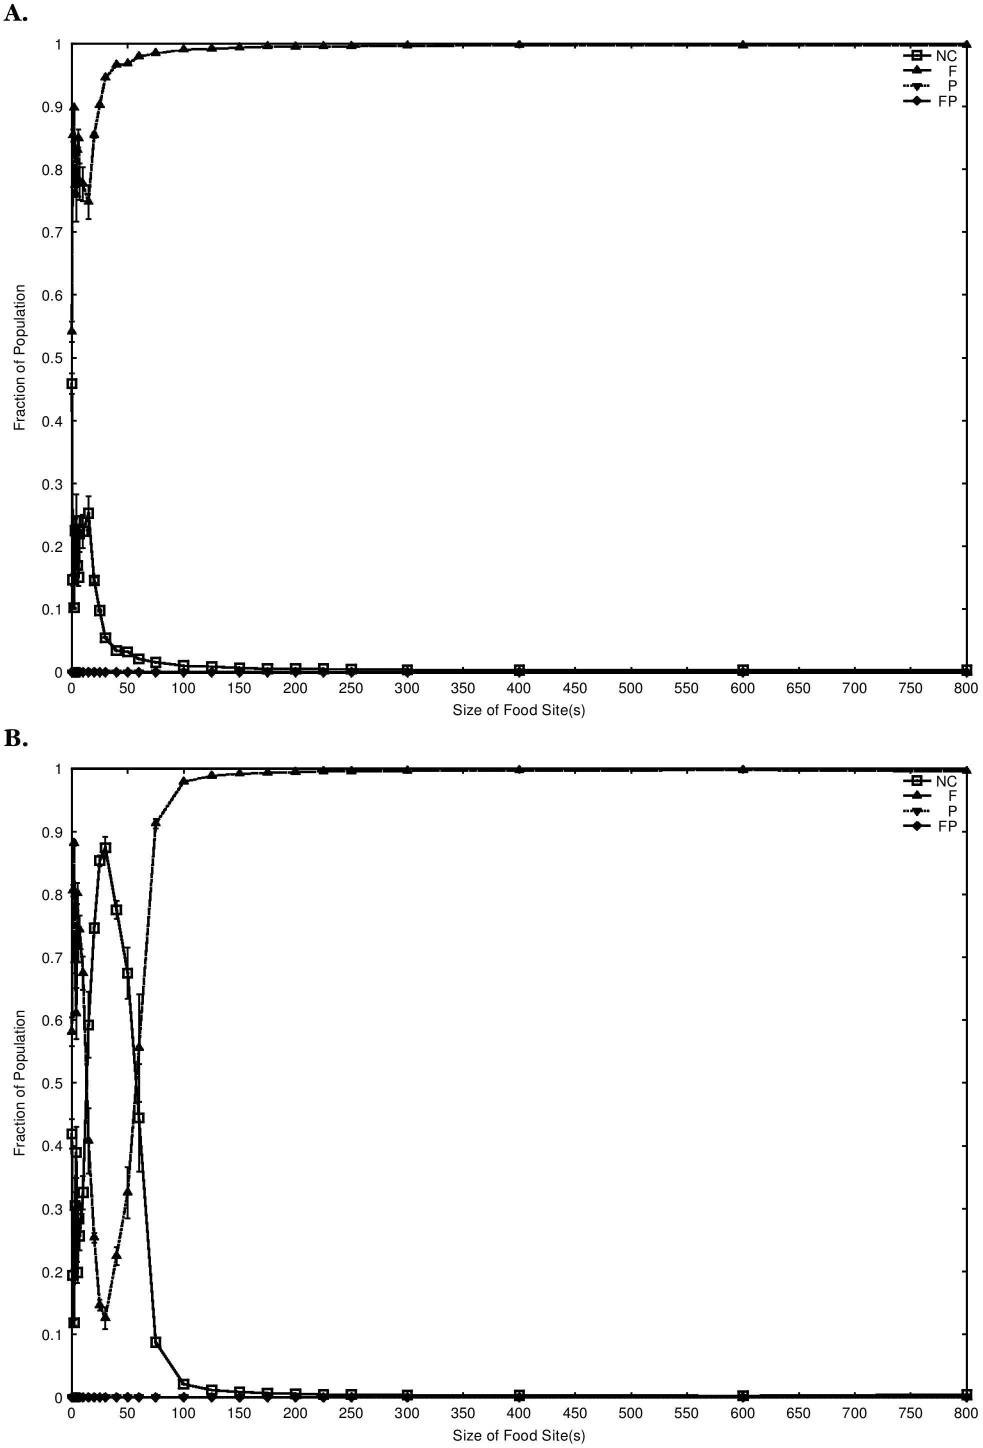

Figure 9. Fraction of agent population with each genotype averaged over last 30,000 time steps of simulations for simulations with different size food sites. Conditions are the same as in Figure 8 except that now spatially constrained rather than random selection and placement is used by the evolutionary process. A) Tournament size of 2. B) Tournament size of 10.

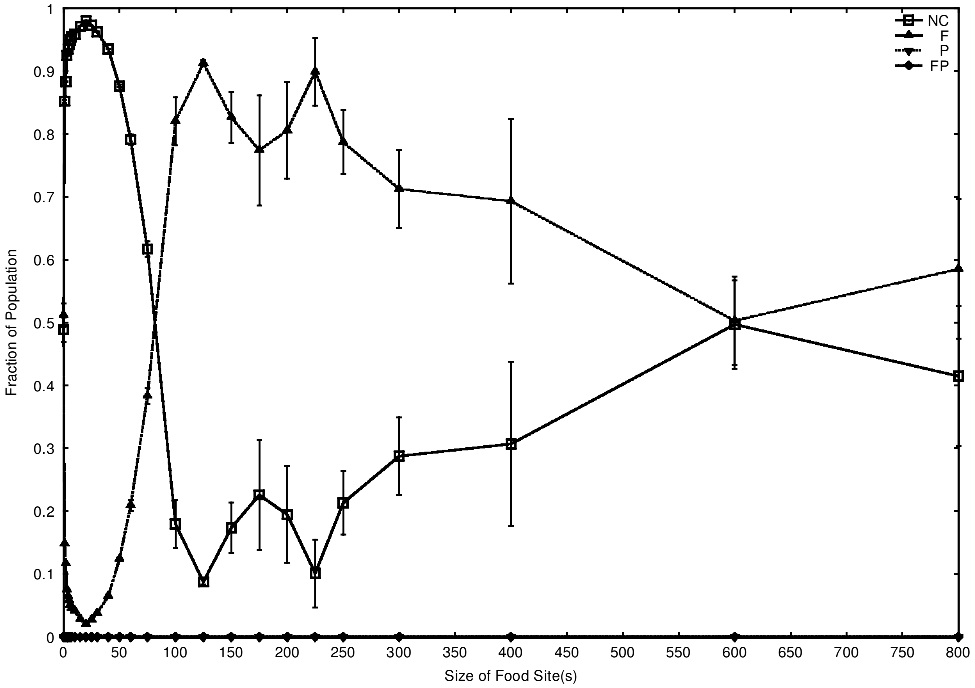

Figure 10. Fraction of agent population with each genotype averaged over last 30,000 time steps of simulations with different size food sites. Same conditions of random selection and placement and a tournament size of 2 as in Figure 8, except that now 56 food sites are present in all simulations. Note the large variance with 100 or more food units per site (this large variance occurs out to 1,600 food units per site).

In contrast to the situation with small food sites, with eight food sites and a large amount of food per site (100 food units or more), a very uniform picture emerged regardless of the evolutionary method used (random vs. spatially constrained selection and placement, either tournament size). In these cases food signaling always disseminated throughout the agent population, with food-communicating agents forming essentially 100% of the population (Figures 8 and 9). This was true up to the largest amounts of food per site examined (1,600 units). This regime of a fairly small number of food sites that were relatively difficult for agents to find, with a large amount of food per site, was the most clearly conducive setting for consistent and complete evolution of food signaling agents.

With 56 food sites, which is roughly 1.5 food sites per 100 cells, results similar to those with 8 food sites were obtained except when a very large amount of food per site was present. Food signaling was generally suppressed when a relatively small amount of food per site was present (< 100 units per site), while it could evolve to be quite prevalent for moderately large food sites (100–300 food units per site). This is illustrated in Figure 10 for random selection and placement and a tournament size of 2. Three regions can be seen in this graph: food signaling suppressed (food per site < 100), food signaling somewhat advantageous (with food per site roughly between 100 and 400), and food signaling essentially neutral (food per site > 500). Although F agents have less overall evolutionary advantage with so many food sites, the pattern seen here is roughly similar to that seen under the same conditions with 8 food sites (Figure 8), except that when more than 500 units of food per site are present there is now no substantial advantage to signaling the discovery of food. With 56 food sites and using random selection again but a tournament size of 10, food signaling never

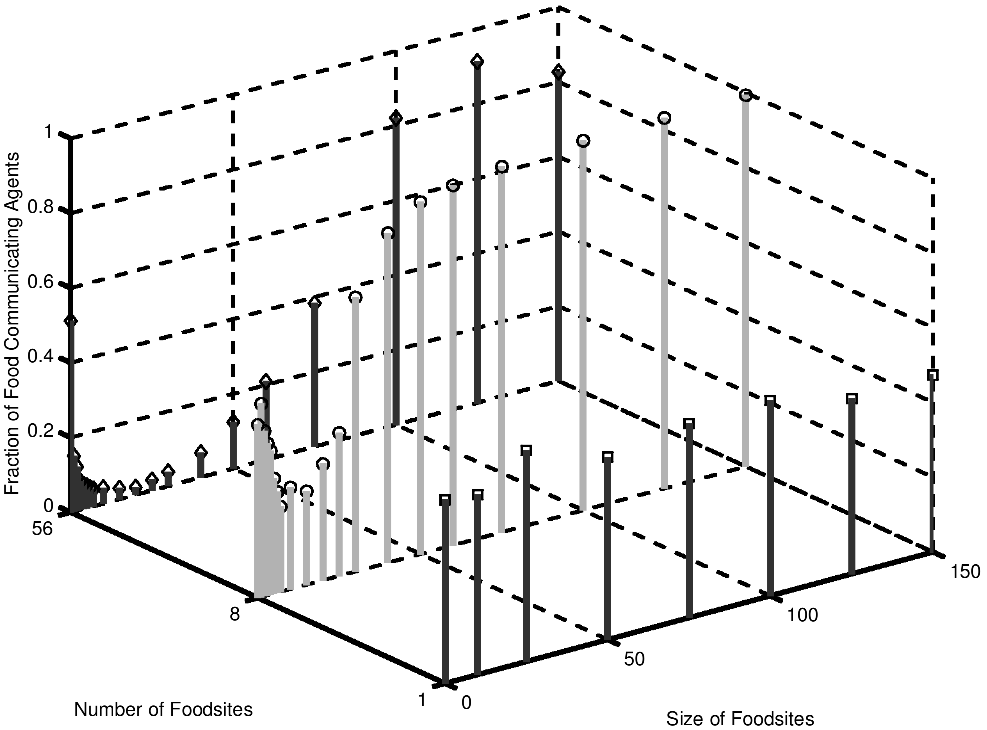

Figure 11. Relative patterns of evolution of food-communicating agents for three different numbers of foodsites as the amount of food per site is varied. Random selection and placement, tournament size 2. Some of the data plotted here overlap with some in Figures 8A and 10, permitting direct comparison.

proved to be advantageous at all. With 56 food sites and spatially constrained selection and placement, the same three-region pattern occurred as is shown in Figure 10. Again, unlike when 8 food sites were present, for the largest food sites (600–1,600 food units per site), food signaling was generally neither advantageous nor disadvantageous. With such very large food sites roughly 50% of the population would ultimately be non-communicating agents and 50% food-signaling agents, regardless of the evolutionary method used, as illustrated by the differences in the rightmost aspects of Figure 10 relative to Figures 8 and 9. The differences between the patterns of evolution of food-communicating agents with varying numbers of food sites are summarized and can be directly compared to each other in Figure 11. As can be seen here, with multiple food sites the general tendency is for increasing site size first to depress food signaling and then to encourage it, but no crisp threshold is evident.

3.6 Impact of Predators on Food Signaling

To assess how the presence of predators affects the evolution of food signaling, all of the simulations described in the preceding section were repeated but now with 20 predators (about 1 predator per 200 cells) present. When a fixed amount of food was present (1,600 units) and a varying number of food sites, the same qualitative patterns as without predators were observed in terms of when food signaling emerged, regardless of selection and placement method or tournament size. For example, with random selection and placement and a tournament size of 2, the same three regions can be identified in a plot of population numbers as the number of food sites is increased (compare Figure 7B with Figure 7A). In the range of 5–15 food sites, approximately

95% of agents communicate about food (90% FP agents + 5% F agents). This is slightly less than the 100% food communication observed for the same number of food sites in the absence of predators, reflecting the fact that signaling about food now has a cost (it can attract predators to the food signaler). Suppression of food communication still occurs when a large number of small food sites are present (e.g., about 10% of agents communicate about food with 64 sites), but this effect requires more and smaller food sites than when no predators are present. In summary, the presence of predators had relatively little impact on the overall pattern of evolution of food signaling under these specific conditions.

The more extensive set of simulations where the number of food sites are held constant and the amount of food per site varies were also all repeated with 20 predators present. With a single food site, virtually no effect was observed on the evolution of food signaling: Roughly 50% of agents evolved to be food signalers, regardless of the selection and placement method, tournament size, or amount of food per site. With 8 food sites and a large amount of food per site, results with predators present were quite similar to when predators were absent, except typically a roughly 5% decrease in food-signaling agents occurred. With 56 food sites and a large amount of food per site, predators inhibited the emergence of food signaling even more, with the fraction of food signalers decreasing about 10% with random selection and placement and 40% with spatially constrained selection and placement, relative to when no predators were present. Finally, with either 8 or 56 food sites and a relatively small amount of food per site, the results were complex and much less predictable: food signalers could decrease or increase compared to when no predators are present, sometimes doing both under the same selection and placement and tournament size scenarios at different amounts of food per site. With small amounts of food per site like this, the already complex interactions influencing emergence of food signaling when predators were absent became even more involved and unpredictable.

For completeness, we note that the varying amounts of food present with 20 predators influenced the evolution of predator signaling in the simulations described in this section in ways consistent with the results described earlier in this article. Recall that with 20 predators and 8 food sites of 200 food units each, we earlier observed that expected predator signaling was *suppressed* with a tournament size of 10 (e.g., Figures 5 and 6B, and associated text). With 8 food sites and varying food per site this was again found to be the case under the more general conditions examined here. Regardless of the amount of food per site and the selection and placement method, more than 90% of agents developed predator signaling when tournament size was 2, while generally less than 10% did so when tournament size was 10.[6] Similar results were obtained with 56 food sites: Regardless of amount of food per site and the selection and placement method, more than 95% of agents developed predator signaling when tournament size was 2, whereas with a tournament size of 10 either virtually no predator signaling emerged (random selection) or at least substantially below baseline levels of predator signaling occurred (spatially constrained selection). Finally, with only a single food site present and a tournament size of 10, suppression of predator signaling did not occur with spatially constrained selection and placement, but some was evident with random selection and placement (roughly 50% of agents evolved to signal the presence of predators in this case). The disappearance or reduction of suppression of predator communication with a tournament size of 10 when just a single food site was present is consistent with the observations made in an earlier section that the extra time spent fleeing by predator-signaling agents lowered their measured fitness due to less time

---

[6] For smaller food sites, especially with spatially constrained selection that encourages the emergence of communication in general, suppression of predator signaling could be much less evident.

spent consuming food. When only one food site is present, extra time fleeing would generally not affect food consumption as most agents do not know where the food site is located.

4 Discussion

In this study we examined through computational modeling the conditions under which initially noncommunicating agents would evolve to warn of nearby predators or to indicate the discovery of food. In these simulations we systematically varied both environmental factors (agent density, predator density, food site density, and amount of food per site) and assumptions about the evolutionary process (role of spatial relationships, dependence of fitness on food acquisition) to examine how these factors influence the appearance and proliferation of communication. Our results confirm previous computational studies showing that it is possible to demonstrate the emergence of signaling between agents in an appropriate setting [1, 2, 12, 29, 30, 44–46]. Further, we have extended these past studies by systematically examining situations in which communication does or does not occur, and by showing that in some contexts communication can be actively prevented/suppressed by the environment when it leads to lower fitness. In the following we first summarize the factors that proved important in our artificial world in determining how signaling evolved, and then discuss some implications for both animal communication and for artificial life models involving communication.

First, we found that there is a threshold density of agents above which food and predator signaling readily appears and proliferates, and below which it does not, even in conducive settings. The threshold is not a sharp cutoff but it is fairly clear-cut (e.g., it occurs with 75 agents, a density of about 2 agents per 100 cells, in Figure 4A). Once agent density is substantially above this value, there is no further benefit of increasing agent density. This threshold can be predicted approximately: It occurs when agent density is sufficiently high that the theoretically expected distance between agents generally put them within communication distance of one another.

Consistent with the above observation, the use of spatially constrained selection and placement by agents during evolution, rather than the random, nonspatial selection and placement used in standard genetic algorithm models, had a moderate but consistent facilitating effect on the evolution of communication. By introducing a tendency to keep communicating agents closer together, for a given total number of agents the *local* density of communicating agents could become higher than the average density, making signaling effective at lower overall agent densities (e.g., Figure 4B). This process was self-reinforcing when food communication was involved because food-signaling agents attracted one another toward food sites. Spatially constrained selection and placement has the additional effect of increasing the frequency with which genetically similar agents interact. Such genetically nonrandom interactions are necessary for kin selection, as has been noted by others [43]. Future simulations in which offspring are placed near agents other than parents could be used to separate the beneficial effects of spatial segregation from kin selection on the evolution of signaling behavior.

In general, it proved easy to evolve predator signaling, requiring only three or four predators in the entire environment (about one predator per 1000 cells) for predator signaling to be advantageous. With small numbers of predators, the density of predators made a difference in the dissemination of predator signaling by agents: As the number of predators present increased, the tendency for agents to use predator signaling increased until almost all agents communicated about predators. This took only about a dozen predators to achieve (Figure 6A). As anticipated, predators also affected emergence of food signaling. While the results could be difficult to predict in situations where a small amount of food per site was present, in other situations predators had a small but persistent tendency to decrease food signaling. This is because the signaling could alert predators to nearby agents that they had not detected by sight.

In many past genetic algorithm and artificial life contexts, a fitness function has been used to guide selection of parents for reproduction [1, 2, 29, 30, 44, 45]. When fitness based upon success in acquiring food was introduced into the evolutionary process (e.g., tournament size 10 as opposed to just basing parent selection upon survival with tournament size 2), occurrence of food signaling occurred much as it did when only survival-based selection was used. However, an unexpected effect of including fitness based upon food consumption was that in some settings, it could inhibit the evolution of signaling the presence of predators (e.g., compare Figure 5A and B). This occurred when predator signaling reduced an agent's risk of death by predation. As explained earlier, the reason this occurred is that the predator-signaling agents had diminished food consumption due to the increased time they spent fleeing. This produced diminished fitness and therefore less probability of selection for reproduction. This effect was evident over a wide range of predator densities (e.g., compare Figure 6A and B). Our conclusion is not so much that one approach to assessing agent fitness (survival alone vs. survival plus food consumption) is preferable, but that these two approaches can produce different outcomes in evolving communication under some conditions. In the future, it would be valuable to examine alternative life histories explicitly, perhaps by changing the age of reproduction or the number of young produced by agents.

Other factors influencing whether or not food signaling specifically evolved to be widespread were the density of food locations and the amount of food at each food site. The setting that was most consistently, convincingly, and categorically conducive to evolving virtually 100% of agents using food signaling was when relatively few food sites were present (e.g., one per 450 cells) but the food at each site was very plentiful. Under these conditions, food was difficult to find without communication, and signaling its presence led other signaling agents to food, increasing their survival and fitness. There was, however, a threshold for this effect in terms of the amount of food per site. With few sites and only a relatively small amount of food per site (roughly 10–50 food units), food signaling was usually strongly suppressed (e.g., Figures 8, 9). This occurred because food-signaling agents attracted each other to the same food site and thus competed for the same limited food, whereas each noncommunicating agent, although less likely to find a food source, was more likely to have it to itself when it did locate food. Finally, food signaling was less likely to emerge when food became much easier to find (e.g., one food site per 60 cells)—even noncommunicating agents could locate food sites and with adequate food per site the environment became essentially neutral with respect to the value of food signaling.

The focus of our simulations was on identifying ecological conditions that favor the evolution of signals for predator alarm or food recruitment. To keep agent behavior simple we assumed that communicating agents always signal when they encounter an appropriate stimulus. Consequently, agents that respond to the calls of others but remain silent when they encounter predators or food do not exist. If signaling provides a direct benefit to the caller, that is, if an agent that calls decreases its own risk of predation, as occurs in meerkats [11], or if an animal that signals the location of food increases the reproduction of close relatives, as occurs in many social insects [5], then this assumption is reasonable since a calling individual will spread its genes faster than a noncalling individual. However, in other cases, such as when there is competition for food with unrelated individuals, an animal that does not advertise a food source but utilizes the food recruitment signals of others should have an advantage over food signalers. The resultant evolutionary outcomes then depend on the behavioral options available. One possibility is that a stable state is reached in which both callers and parasitic scroungers persist [8] or, if there are spatial constraints influencing interactions, more dynamic outcomes are possible, including persistent cooperation [34]. Alternatively, individuals may adopt more complicated conditional strategies. For example, when the relative benefits and costs of signaling fit a Prisoner's Dilemma, and there is a high probability of future interaction, "Tit-for-Tat" (TFT), in which one player always initiates an interaction with a partner by cooperating but on subsequent occasions copies its partner's previous behavior, can outcompete exploitative strategies [3] and provide a transitional step toward other strategies including unconditional cooperation [33]. TFT and many other conditional strategies require the ability to recognize individuals and remember their past behavior. Thus, full exploration of the consequences of signal exploitation must await more sophisticated agents that have the ability to modify their behavior in response to interactions with other specific individual agents.

The implications of this study are twofold. First, in terms of animal communication, these modeling results make predictions about when animal signaling would be expected to evolve. For example, the ease with which warning signals about predators evolved to be widespread in the presence of just a few predators suggests that predator signaling should be a common occurrence in nature. In accord with this prediction, most social vertebrate species give at least one type of signal when an attacking predator is detected [5, 28]. Evidence from ground squirrels [39] and dwarf mongoose [36] indicates that such signals are costly because they enhance risk of predation to the caller. These costs are likely outweighed by indirect or direct benefits that accrue to callers either because receivers are related to callers [39, 40] or because the resultant behavior of the recipient decreases its risk of predation. The greater difficulty encountered in evolving predator communication in the model when current food stores were made the basis of selection for reproduction reflects the fact that, in reality (unlike the model), animals use various strategies, such as grouping, to reduce predation without impairing food consumption. There are a number of examples of animals preferring to feed in suboptimal habitats to avoid predation [23, 31, 47]. This leads to reduced growth rates and probably reduced survival. Presumably the animals are still behaving in an optimal way to balance the increased risk of predation against the decreased food availability.

Food signaling evolved most readily when a few rich food sites existed that were difficult to find. In other situations, signaling about food resulted in multiple costs to the signaler, that is, food intake was reduced by recruitment of competitors and predation risk was increased by attracting predators. In addition, food calling also appeared more readily when callers aggregated presumably because the benefits of finding food through repeated interaction exceeded the occasional costs of signaling. Similar predictions have been obtained from game-theory models of this behavior [32]. Not surprisingly, food signals are uncommonly reported for birds and mammals and are invariably associated with the discovery of rich patches of difficult-to-find food [6, 7, 9, 13, 48, 49]. In addition, other benefits to calling can often be ascribed to the signaler. For example, ravens that call at carcasses attract vagrants that help to ward off territorial pairs [21]. Calling at food increases group size and enhances detection of approaching predators in house sparrows [14] or facilitates tracking of large ephemeral insect swarms by cliff swallows [6]. Rhesus macaque females avoid punishment by males if they call when food is disovered [20]. Finally, naked mole rats increase the survival of close genetic relatives by vocalizing and providing scent trails to food discoveries [25]. Thus, food calling is most likely when either the cost of sharing is low or attracting conspecifics provides some independent benefit.

The second implication of our results is that future artificial life research on the evolution of communication must give careful attention to the assumptions made about the evolutionary process and the spatial relationships of agents. Such assumptions can have a substantial impact on the evolution of communication in sometimes surprising ways. Incorporating spatial relationships, something not addressed by traditional genetic algorithms nor in some past studies of evolving communication, proved to be very important. In part, this reflects the fact that signaling inherently involves multiple cooperating and interacting agents, unlike in a number of other past artificial life studies where agents are isolated and do not interact with each other. In some cases, results found when spatial relations were considered in the selection of parents and placement of new agents made a qualitative difference in the results. Sometimes the attraction of communicating but not of noncommunicating agents to one another, such as at food sites, essentially provided positive feedback that facilitated the evolution of signaling. Such issues, as well as numerous potential variations to agent behavior and parameter settings (e.g., food reserves at birth; ignoring predator signals from distant agents if starving; minimum age for reproduction), surely merit further study.

Acknowledgments

We thank the anonymous reviewers for their helpful comments. This work was supported in part by the University of Maryland Center for Neuroscience.

References

1. Ackley, D., & Littman, M. (1994). Altruism in the evolution of communication. In R. Brooks & P. Maes (Eds.), *Artificial Life IV* (pp. 40–48). Cambridge, MA: MIT Press.
2. Arita, T., & Koyama, Y. (1998). Evolution of linguistic diversity in a simple communication system. *Artificial Life, 4,* 109–124.
3. Axelrod, R., & Hamilton, W. (1981). The evolution of cooperation, *Science, 211,* 1390–1396.
4. Benz, J. (1993). Food-elicited vocalizations in golden lion tamarins: Design features for representational communication. *Animal Behavior, 45,* 443–455.
5. Bradbury, J. W., & Vehrencamp, S. L. (1998). *Principles of animal communication.* Sunderland, MA: Sinauer.
6. Brown, C. R., Brown, M. B., & Shaffer, M. L. (1991). Food-sharing signals among socially foraging cliff swallows. *Animal Behavior, 42,* 551–564.
7. Caine, N. (1995). Factors affecting the rates of food calls given by red-bellied tamarins. *Animal Behavior, 50,* 53–60.
8. Caraco, T., & Giraldeau, L.-A. (1991). Social foraging: Producing and scrounging in a stochastic environment. *Journal of Theoretical Biology, 153,* 559–583.
9. Chapman, C. A., & Lefebvre, L. (1990). Manipulating foraging group size: Spider monkey food calls at fruiting trees. *Animal Behavior, 39,* 891–896.
10. Cheney, D., & Seyfarth, R. (1990). *How monkeys see the world.* Chicago: Chicago University Press.
11. Clutton-Brock, T., O'Riain, M., Brotherton, P., Gaynor, D., Kansky, R., Griffin, A., & Manser, M. (1991). Selfish sentinels in cooperative mammals. *Science, 284,* 1640–1644.
12. de Bourcier, P., & Wheeler, M. (1995). Aggressive signaling meets adaptive receiving. In F. Moran, A. Moreno, J. Merelo, & P. Chacon (Eds.), *Advances in artificial life* (pp. 760–771). Berlin: Springer.
13. Dittus, W. P. J. (1984). Toque macaque food calls: Semantic communication concerning food distribution in the environment. *Animal Behavior 32,* 470–477.
14. Elgar, M. A. (1986). House sparrows establish foraging flocks by giving chirrup calls if the resources are divisible. *Animal Behavior 34,* 169–174.
15. Elowson, A., Tannenbaum, P., & Snowdon, C. (1991). Food-associated calls correlate with food preferences in cotton-top tamarins. *Animal Behavior 42,* 931–937.

16. Hamilton, W. (1996). The genetical evolution of social behavior. *Journal of Theoretical Biology 7*, 1–51.

17. Hauser, M. (1996). *The evolution of communication.* Cambridge, MA: MIT Press.

18. Hauser, M., & Marler, P. (1993). Food-associated calls in rhesus macaques (*Macaca mulatta*): I. Socioecological factors. *Behavioral Ecology 4*, 194–205.

19. Hauser, M., & Marler, P. (1993). Food-associated calls in rhesus macaques (*Macaca mulatta*): Costs and benefits of call production and suppression. *Behavioral Ecology 4*, 206–212.

20. Hauser, M. D., Teixidor, P., Field, L., & Flaherty, R. (1993). Food-elicited calls in chimpanzees: Effects of food quantity and divisibility. *Animal Behavior 45*, 817–819.

21. Heinrich, B., & Marzluff, J. M. (1991). Do common ravens yell because they want to attract others? *Behavioral Ecology and Sociobiology, 28*, 13–21.

22. Holland, J. (1975). *Adaptation in natural and artificial systems.* Ann Arbor: University of Michigan Press.

23. Holmes, W. (1984). Predation risk and foraging behavior of the hoary marmot in Alaska. *Behavioral Ecology and Sociobiology, 15*, 293–302.

24. Hutchins, E., & Hazlehurst, B. (1995). How to invent a lexicon. In N. Gilbert & R. Conte (Eds.), *Artificial societies* (pp. 157–189). London: UCL Press.

25. Judd, T., & Sherman, P. (1996). Naked mole-rats recruit colony mates to food sources. *Animal Behavior, 52*, 957–969.

26. Koza, J. (1992). *Genetic programming.* Cambridge, MA: MIT Press.

27. Luke, S., & Spector, L. (1998). A revised comparison of crossover and mutation in genetic programming. In J. Koza, W. Banzhaf, K. Chellapilla, K. Deb, M. Dorigo, D. Fogel, M. Garzon, D. Goldberg, H. Iba, & R. Riolo (Eds.), *Proceedings of the Third Annual Conference on Genetic Programming* (pp. 208–214). San Francisco: Morgan Kaufmann.

28. Macedonia, J. M., & Evans, C. S. (1993). Variation among mammalian alarm call systems and the problem of meaning in animal signals. *Ethology, 93*, 177–197.

29. MacLennan, B. (1990). *Evolution of communication in a population of simple machines.* (Tech. Rep. No. CS-90-99). Knoxville: University of Tennessee.

30. MacLennon, B. (1991). Synthetic ethology: An approach to the study of communication. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmusen (Eds.), *Artificial life II* (pp. 631–658). Redwood City, CA: Addison-Wesley.

31. Milinski, M., & Heller, R. (1978). Influence of a predator on the optimal foraging behaviour of sticklebacks (*Gasterosteus Aculeatus*). *Nature, 275*, 642–644.

32. Newman, J. A., & Caraco, T. (1989). Cooperative and non-cooperative bases of food-calling. *Journal of Theoretical Biology, 141*, 197–209.

33. Nowak, M., & Sigmund, K. (1993). A Strategy of win stay, lose shift that outperforms tit-for-tat in the prisoners-dilemma game. *Nature, 364*, 56–58.

34. Nowak, M., & May, R. (1992). Evolutionary games and spatial chaos. *Nature, 359*, 826–829.

35. Pereira, M., & Macedonia, J. (1991). Ringtailed lemur antipredator calls denote predator class, not response urgency. *Animal Behavior 41*, 543–544.

36. Rasa, O. A. (1989). The costs and effectiveness of vigilance behavior in the dwarf mongoose: Implications for fitness and optimal group size. *Ethology, Ecology and Evolution, 1*, 265–282.

37. Reggia, J., Schulz, R., Uriagereka, J., & Wilkinson, J. (2000). *A simulation environment for evolving multi-agent communication.* Tech. Rep. College Park: University of Maryland, Dept. of Computer Science.

38. Seyfarth, R., Cheney, D., & Marler, P. (1980). Vervet monkey alarm calls: Semantic communication in a free-ranging primate. *Animal Behavior 28*, 1070–1094.

39. Sherman, P. W. (1977). Nepotism and the evolution of alarm calls. *Science, 197*, 1246–1253.

40. Sherman, P. (1985). Alarm calls of Belding's ground squirrels to aerial predators: Nepotism or self-preservation? *Behavioral Ecology and Sociobiology 17*, 313–323.

41. Steels, L., & Kaplan. F. (1998). Stochasticity as a source of innovation in language games. In C. Adams, R. Belew, H. Kitano, & C. Taylor (Eds.), *Artificial Life VI* (pp. 368–376). Cambridge, MA: MIT Press.

42. Steels, L., & Kaplan, F. (1999). Collective learning and semiotic dynamics. In D. Floreano, J. Nicoud, & F. Mondada (Eds.), *Advances in artificial life* (pp. 679–688). Berlin: Springer.

43. Wade, M. (1980). Kin selection: Its components. *Science, 210*, 665–667.

44. Wagner, K. (1999). Habitat, communication and cooperative strategies. In W. Banzhof, J. Daida, A. Eiben, M. Garzon, V. Honover, M. Jakiela, & R. Smith *Proceedings of the Genetic and Evolutionary Computation Conference* (pp. 694–701). San Francisco: Morgan Kaufmann.

45. Wagner, K. (2000). Cooperative strategies and the evolution of communication. *Artificial Life, 6*, 149–179.

46. Werner, G., & Dyer, M. (1991). Evolution of communication in artificial organisms. In C. G. Langton, C. Taylor, J. D. Farmer, & S. Rasmusen (Eds.), *Artificial life II* (pp. 659–687). Redwood City, CA: Addison-Wesley.

47. Werner, E., Gilliam, J., Hall, D., & Mittelbach, G. (1983). An experimental test of the effects of predation risk on habitat use in fish. *Ecology, 64*, 1540–1548.

48. Wilkinson, G. S., & Boughman, J. W. (1998). Social calls coordinate foraging by greater spear-nosed bats. *Animal Behavior 55*, 337–350.

49. Wrangham, R. (1977). Feeding behavior of chimpanzees in Gombe National Park, Tanzania. In T. Clutton-Brock (Ed.), *Primate ecology* (pp. 503–538). London: Academic Press.