Beating your Neighbor to the Berry Patch

Alan R. Rogers*

November 12, 2020

Abstract

When similar foragers compete for a resource that ripens (or improves in some other way) with time, no strategy (pure or mixed) can be evolutionarily stable. There is a unique mixed Nash Equilibrium (NE), which is stable against pure-strategy invaders provided that either the cost of visiting the resource or number of competitors is large. This equilibrium, however, is never stable against invading mixed strategies that are sufficiently similar to the NE.

In computer simulations, this mixed-strategy instability was not observed. The process converged to the neighborhood of the NE whenever the parameters implied that the NE was stable against pure-strategy invaders.

The mixed-strategy instability was not observed either in experiments with human subjects. After an initial period of familiarization, the behavior of human subjects was close to that predicted by the NE whenever the model predicted stability against pure-strategy invaders.

Key words: foraging, evolutionary game theory, competition, ESS, auction

1 Introduction

Every summer, my backyard witnesses a conflict between humans and birds, all of whom wish to eat the same strawberries. Those who wait until the berries are ripe eat none, for by then the others have come and gone. All of us eat sour berries or none at all, and none of us are happy about it.

Such interactions must be common in nature. They occur whenever

1. Several individuals compete for the same resource.
2. The resource improves in value over time.
3. Some cost is involved in attempting to harvest the resource whether one succeeds in harvesting it or not.

*Dept. of Anthropology, Univ. of Utah, S.L.C. UT 84112, USA.

4. Harvesting the resource ruins it for those who come later.

I know of three examples from ethnography. (1) The Barí are a horticultural people in Venezuela whose fishing methods have been described by Bennett [1]. The Barí often fish by dumping poison into pools, where the stream is deep and slow. This kills most or all of the fish. New fish enter the pool only slowly, so that the pool improves in value over a period of weeks. The Barí defend territories, but different villages nonetheless exploit the same pools. The Barí therefore face a dilemma. If they wait until the pool is full of fish, another village may exploit the pool first. But if they go too soon, the pool is hardly worth exploiting. (2) The Hadza are a foraging people in Tanzania [2]. Like most tropical foragers, the Hadza enjoy honey. Hives improve in quality during the spring and summer, so it is best not to exploit them too soon. But the Hadza also compete for honey not only with each other but also with birds and badgers. Humans and badgers both destroy hives when they exploit them, so little is left for subsequent foragers. (3) The Aché are a population in Paraguay whose economy involves both foraging and gardening. Kim Hill (personal communication) tells me that garden products are seldom allowed to ripen because children roam the gardens foraging for themselves. Parents who waited for the produce to ripen would harvest nothing. Hill has worked with the Aché for decades but has yet to eat a ripe watermelon.

These examples show that the interaction in my back yard is not an isolated example. It illustrates a problem that must have confronted our ancestors for a very long time. Thus, it makes sense to ask what strategy would have been favored by natural selection. Below, I introduce a model that answers this question. First, however, I motivate the theory by showing how real people respond to similar dilemmas in classroom experiments.

2 A classroom experiment

Subjects interacted with each other via a computer program, which provided instructions, calculated scores, and kept track of each subject's choices. The screen is shown in figure 1.

Subjects play in groups of five. In each round of the game, each subject chooses between "going fishing," which yields a certain return of 2 lab dollars, and attempting to harvest the berrypatch. Those who attempt the berrypatch choose a value at which to harvest. A subject who chooses the value v will gain:

v lab dollars	if no other subject chooses a value as small
v/n lab dollars	if n subjects tie for the smallest value
nothing	if some other subject chooses a smaller value

At the end of the game, subjects are paid 0.03 US dollars for each lab dollar.

Figure 2 shows the results from two experiments, each with 5 subjects, and totalling 172 trials. The students ignored the berry patch about half of the time. On those occasions

Figure 1: Computer screen used in berrypatch game

```
 ------------------------------------- -------------------------------------
|                                     ||last play (Fish) earned $2          |
|  INSTRUCTIONS                       ||assets = $2                         |
|A small berry patch ripens during the ||opponents = 4                      |
|summer.  It is worthless in April and ||round: 2 of 100                    |
|worth $10 in August if the berries   ||                                    |
|are still there.  The first visitor  ||                                    |
|harvests all the berries; later      ||                                    |
|visitors get nothing.  In case of a  ||                                    |
|tie, the patch is divided among the  ||                                    |
|winners.                             ||                                    |
|                                     ||                                    |
|If you don't visit the berrypatch at | -------------------------------------
|all, you can go fishing, which always | -------------------------------------
|yields $2.                           ||Choices:                            |
|                                     ||                 f: go fishing      |
|You are competing against 4 other    ||value btw 0 and 10:                 |
|foragers. The game will last 100     ||  visit patch when it reaches value |
|rounds.                              ||  (value need not be an integer)    |
|                                     | -------------------------------------
|At the end, I will pay you 0.03 real | -------------------------------------
|dollars for every "lab dollar".      ||Your choice: _                      |
|                                     ||                                    |
 ------------------------------------- -------------------------------------
```

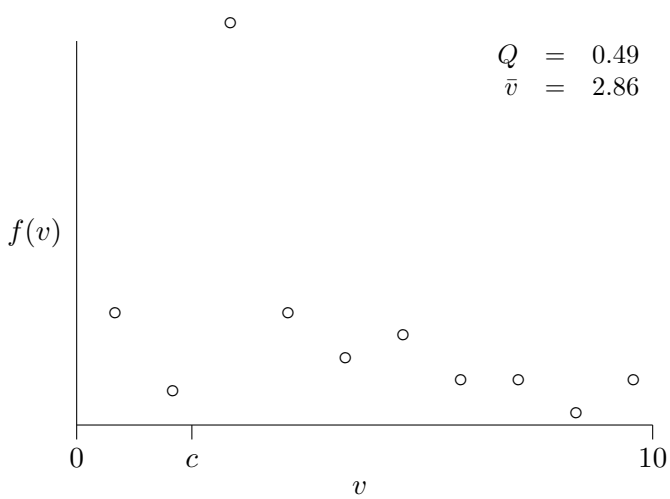

Figure 2: Experimental Results

Data from two experiments, each with 5 subjects and totaling 172 trials. Q is the frequency with which "go fishing" was played, and $\bar{v}$ is the mean strategy among the other plays.

when they did visit it, they were most likely to visit when the patch's value barely exceeded the opportunity cost (the payoff from going fishing).

Now these students are not foragers, but each of them is descended from a long line of foragers. It seems possible that our species has evolved a brain that is well equipped to find the optimal solution to such problems. To find out whether these students reached an optimal solution, we need a model.

3 Model

In the experiment, the berrypatch game was played 100 times. The model, however, will deal only with a single round of the game: with what is called the "stage game." There is an easy way to justify this simplification: Because the game is played a fixed number of times, the Nash Equilibrium (NE) of the repeated game must involve repetition of the ESS of the stage game [10, pp. 155–156]. This easy justification is suspect, however, because it depends on a feature of the game that is unrealistic. In nature animals do not compete against an unchanging group of competitors for a fixed number of rounds. In some ecological contexts, the game is played only once; in others it is played an unpredictable number of times against a variable group of competitors. Games that are repeated a fixed number of times are evolutionary novelties, and we may have evolved no adaptation to them. Nonetheless, it is always best to start simple, so this paper will deal with a single repetition of the stage game. The results will be relevant to games that are played only once, but should be applied only with caution to repeated games.

In the model, $K + 1$ foragers each want a single resource, but only one forager—the one who first attempts to harvest it—will have it. The value v of the resource increases from 0 at the beginning of the season to a maximum of V as the resource ripens. Foragers who attempt to harvest the resource must decide how ripe to let the resource get before attempting to harvest it. The first forager to visit the resource gets it and thus gains v, its value when harvested. Those who try to harvest the resource later get nothing. If n individuals arrive at the same time, then each has an equal probability of success so that the expected payoff is v/n. Those who ignore the resource altogether can engage in some other activity that yields a certain payoff of c. I will refer to this alternative activity as "going fishing." For notational simplicity, I set $V = 1$, which amounts to measuring all benefits and costs as proportions of V, the maximum potential benefit.

T. Bergstrom (personal communication) observes that this can also be interpreted as a model of an auction with $K+1$ competitors. Each competitor first decides whether to pay an entry fee c, which allows participation in the auction. Participants then choose a bid, $b := V - v$, and the prize goes to the highest bidder. In the literature on auctions, most authors have been concerned either with the case in which each participant values the prize differently but knows only his own valuation [13] or with the case in which each participant has a private estimate of the prize's unknown value [14]. I take a different approach here,

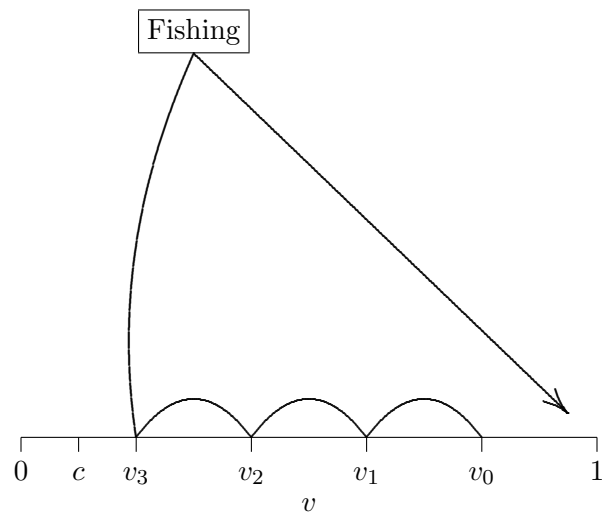

Figure 3: Dynamics with pure strategies

assuming initially that the value of the resource is known with certainty and is the same for all competitors. Later, I relax this assumption using computer simulations.

These interactions are assumed to take place within some large population. Each generation, the members of the population are randomly divided into groups of size $K+1$, and each group then plays the berrypatch game. In evolutionary game theory, we are interested not in the payoff to some strategy within a particular group, but in the average payoff to that strategy across the population.

3.1 Pure strategies

An *evolutionarily stable strategy* (ESS) is a strategy that resists invasion by all alternative strategies [9]. In the present model, no pure strategy can be an ESS. To see why, first note that it never pays to choose $v < c$ because one can always do better than this by going fishing. Thus, I ignore values of v that are less than c. Next, suppose that nearly all of the population plays the strategy labeled v_0 in figure 3. Since everyone in each group is playing the same strategy, the benefits are divided $K+1$ ways and each individual earns $v_0/(K+1)$. A rare mutant who played $v = v_1$ (where $v_1 < v_0$) would always beat its neighbors to the berry patch. When rare, the mutant almost always occurs in groups by itself and therefore does not have to share the resource. Consequently, it will earn v_1. It will increase in frequency when rare provided that $v_1 > v_0/(K+1)$. But this is always true if v_1 is sufficiently close to v_0. Thus, any pure strategy between $v = c$ and $v = 1$ can be invaded by mutants playing a slightly smaller value of v. No such strategy can be an ESS. The strategy $v = c$ is not an ESS either, because each member of a population

playing this would earn $c/(K+1)$ and could do better by going fishing. The only other pure strategy is "go fishing," which earns a payoff of c. But fishing is not an ESS either, for a population of fishers could be invaded by mutants playing $v > c$. Thus, there are no symmetrical equilibria in pure strategies.

There may be asymmetrical Nash equilibria, but I doubt it. For example, suppose that within each group individual A plays $v = c$ and the rest go fishing. Then each individual earns c, yet this is not a Nash equilibrium because A could do better by playing some v greater than c. Even if this were a Nash equilibrium, it would pose a coordination problem. Asymmetrical equilibria are feasible only if there is some means of deciding in advance which player will play which role. Thus, it is of interest to consider the case in which asymmetrical equilibria are impossible, and there are no pure-strategy equilibria at all.

If there is no equilibrium in pure strategies, how can we expect the population to behave? One possibility is that the dynamics will be cyclical as suggested in figure 3. There, the population is initially fixed at $v = v_0$, but the value of v gradually declines as successively smaller mutants invade. A mutant with strategy v receives payoff v when rare and $v/(K+1)$ when common. Eventually, v falls so low that this latter payoff is less than c, the payoff from going fishing. Fishing therefore increases in frequency, as indicated by the path from v_3 to "fishing" on the figure. This increase continues until the resource is rarely harvested, and mutants playing large v can invade. This brings us back to our starting point. This story suggests that foraging behavior might exhibit cyclical dynamics, a point to which we will return.

3.2 A mixed equilibrium

But there is another possibility: foragers may randomize their strategies by choosing v from a probability distribution. Let I denote a strategy that chooses value $v = x$ with probability density $f(x)$ and chooses not to visit the resource at all with probability Q. I assume that all values chosen by I fall within an interval $[L, U]$, where $c \leq L < U \leq 1$. In other words, L is the lowermost value ever chosen and U the uppermost.

The value of L is easy to determine. Suppose that $L > c$. Then a population playing I could be invaded by a mutant playing a fixed value of v that lay between c and L. Thus, I can be an ESS only if $L = c$. Analogous arguments show that $f(x) > 0$ over the entire interval $c < x < U$ and that f must be a "pure density": it cannot give non-zero probality to any point in this interval. Before determining the value of U, I must first derive some formulas.

The probability density $f(v)$ is closely related to two other functions. The *survival function* $s(v)$ is the probability that a forager playing strategy I will not have visited the resource by the time its value is v. It equals

$$s(v) = \int_v^U f(x)dx + Q \tag{1}$$

The *hazard function* $h(v)$ is the conditional probability density of a visit when the resource has value v, given that no visit has yet been made. For convenience, I record here a series of relationships among these functions, which are well-known within demography and survival analysis [7, p. 6].

$$\begin{aligned} s(v) &= \exp\left[-\int_c^v h(x)dx\right] && (2) \\ f(v) &= -ds(v)/dv && (3) \\ &= h(v)s(v) && (4) \\ h(v) &= -d\ln s(v)/dv && (5) \end{aligned}$$

I denote by $\Pi(v, I^K)$ the payoff to a forager playing pure strategy v against K opponents playing I. This payoff equals

$$\Pi(v, I^K) = vs(v)^K \tag{6}$$

since $s(v)^K$ is the probability that none of the K opponents visit the resource by the time its value is v. The payoff to the mixed strategy I is an expected value:

$$\Pi(I, I^K) = \int_c^U f(v)\Pi(v, I^K)dv + Qc \tag{7}$$

To find a formula for s, I make use of the fact that if I is an ESS, then

$$\Pi(x, I^K) = \Pi(y, I^K) \tag{8}$$

for any two pure strategies x and y that are played by I with positive probability [9, pp. 182–183]. In other words, all strategies receive the same payoff when playing against I. Consequently, the graph of $\Pi(v, I^K)$ against v must be flat, and $d\Pi(v, I^K)/dv = 0$. This implies that

$$-\frac{ds(v)}{dv} = \frac{s(v)}{Kv}$$

The left-hand side of this expression above equals $f(v)$ (see equation 3). Consequently, the right-hand side must equal $h(v)s(v)$ (see equation 4). The hazard function is therefore

$$h(v) = 1/(Kv) \tag{9}$$

Substituting into equations 2 and 4 gives the survival and density functions:

$$\begin{aligned} s(v) &= (c/v)^{1/K} && (10) \\ f(v) &= \frac{c^{1/K}}{Kv^{1+1/K}} && (11) \end{aligned}$$

A forager will fail to visit the resource with probability

$$Q := s(U) = (c/U)^{1/K}$$

I assume that $c > 0$: that fishing pays, so that visiting the resource entails an opportunity cost. This requires that $Q > 0$—that foragers go fishing at least occasionally. Since fishing is part of the mixed equilibrium, the return from attempting to harvest the berry patch must on average equal the return c from fishing. On average, therefore, the net benefit from foraging must equal c, the opportunity cost. This insight can be verified by substituting equation 10 into equation 6, which yields $\Pi(v, I^K) = c$ irrespective of v. Since each pure strategy yields payoff c against I, it follows that an individual playing I against I will earn c too.

We are now in a position to determine the value of U, the uppermost pure strategy that is ever played by I. Consider the fate of a rare mutant playing $v = 1$ against a population playing I. The mutant's fitness is

$$\Pi(1, I^K) = s(U)^K = c/U$$

Meanwhile, the I-strategists each earn $\Pi(I, I^K) = c$. The mutant's fitness is greater unless $U = 1$. Thus, U must equal 1 if I is an ESS. Strategy I chooses pure strategies from the entire interval between c and 1; it ignores the resource with probability

$$Q = s(1) = c^{1/K} \tag{12}$$

This result apparently holds in contexts more general than the present model, for it has also been derived in related models of auctions with entry fees [8, Eqn. 9]. The mean value of v among foragers who visit the resource is

$$\bar{v} = (1 - Q)^{-1} \int_c^1 v f(v) dv = \begin{cases} -\frac{c \ln c}{1-c} & \text{if } K = 1 \\ \frac{c^{1/K} - c}{(K-1)(1-c^{1/K})} & \text{otherwise} \end{cases} \tag{13}$$

3.3 Stability of the mixed equilibrium

These results guarantee that I is a Nash equilibrium but not that it is evolutionarily stable. When everyone plays I, all strategies receive equal payoffs. I resists invasion only if the fitness of any alternative strategy would decline as its frequency increased. To express the condition under which this is true, we need notation for payoffs against a heterogeneous mixture of opponents. Let $\Pi(x, J^1 I^{K-1})$ denote the payoff to some strategy x against K opponents of whom 1 plays strategy J and $(K - 1)$ play I. In the appendix, I show that I resists all invasions if

$$\Pi(I, J^1 I^{K-1}) > \Pi(J, J^1 I^{K-1}) \tag{14}$$

for all (pure or mixed) strategies J that differ from I. I resists invasion by all pure strategies when

$$Q > 1/2, \quad \text{or equivalently,} \quad 2^K c > 1. \tag{15}$$

The NE resists pure-strategy invaders only if foragers are more likely to ignore the berry-patch than to try to harvest it. The larger the values of c and K, the more likely this is

to be so. This condition refers only to invasion by pure-strategy mutants. Mixed-strategy invaders are a separate issue. In the appendix, I show that the NE is *never* stable against mixed-strategy invaders that are sufficiently similar to it. The NE is therefore not an ESS.

Before concluding that this game has no equilibria, however, we must consider the possibility that the equilibrium is polymorphic, with several strategies maintained at stable frequencies. Maynard Smith [9, appendix D] discusses the conditions under which mixed equilibria can be construed to represent mixtures of individuals each of whom plays a pure strategy (see also [12, 15, 16]). I have not, however, found any treatment of mixtures of mixed strategies. I therefore attack the problem myself in appendix B.

There, I consider a polymorphic population of individuals playing various pure or mixed strategies. It turns out that the necessary and sufficient conditions for evolutionarily stability in this polymorphic population are identical to those for a mixed ESS in a homogeneous populations: conditions 8 and 14 must both hold. It follows that all of the results derived above apply also when I is interpreted as a mixture of individuals of different kinds. For example, $f(v)$ (defined in equation 11) describes the probability that a random individual drawn from this mixed population will play pure strategy v. However, since inequality 14 does not hold, no such polymorphic population can be evolutionarily stable.

What possibilities remain? The dynamics may be chaotic or cyclical. Furthermore, a strategy may be dynamically stable yet fail to be an ESS [5, p. 203].

4 Computer simulations

4.1 Mixtures of pure strategies

I consider first a simplified world in which no individual plays a mixed strategy. Each simulation begins with all individuals playing $v = 1$. The first event in the life cycle is mutation, which assigns new strategies to one per cent of the population. Of these mutants, half become fishers and half are assigned a value of v chosen at random on the interval between 0 and 1. After mutation, fitnesses are assigned using the model above. Reproduction is haploid, with each individual producing offspring in proportion to her fitness.

Figures 4–5 show simulations with increasing values of $2^K c$. Inequality 15 indicates that these should show increasing stability, and indeed they do. In the less stable simulation (figure 4), $2^K c$ is well below unity, so inequality 15 implies that I can be invaded. As the figure shows, $\bar{v}$ and Q both oscillate wildly. In figure 5, $2^K c$ exceeds unity, so inequality 15 tells us that I cannot be invaded by pure strategies. In this case, $\bar{v}$ and $\bar{Q}$ each converge rapidly toward the unstable equilibrium described by equations 13 and 12.

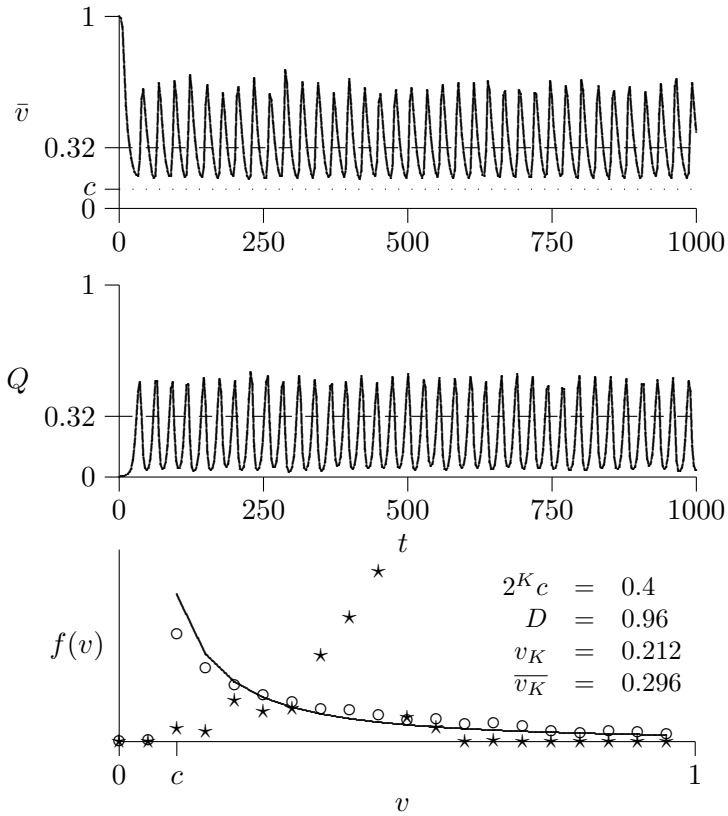

Figure 4: Simulation in which $c = 0.1$ and $K = 2$

The simulation assumes that $c = 0.1$ and $K = 1$. In each generation there were 3333 groups of size 3. In the upper two panels the dashed lines show $\bar{v} = 0.25$ (upper panel) and $Q = 0.1$ (middle panel), the values at the unstable equilibrium as given by equations 13 and 12. In the lower panel, the stars and circles show empirical frequency distributions of the strategy variable v. The distribution shown with stars was calculated from the simulation's final generation, while that shown with circles aggregates over a large number of generations—all generations since the first in which $\bar{v}$ fell to the value predicted by equation 13. The former distribution summarizes 9653 individuals, the latter 7,470,959. The solid line shows the distribution predicted at the (unstable) mixed equilibrium.

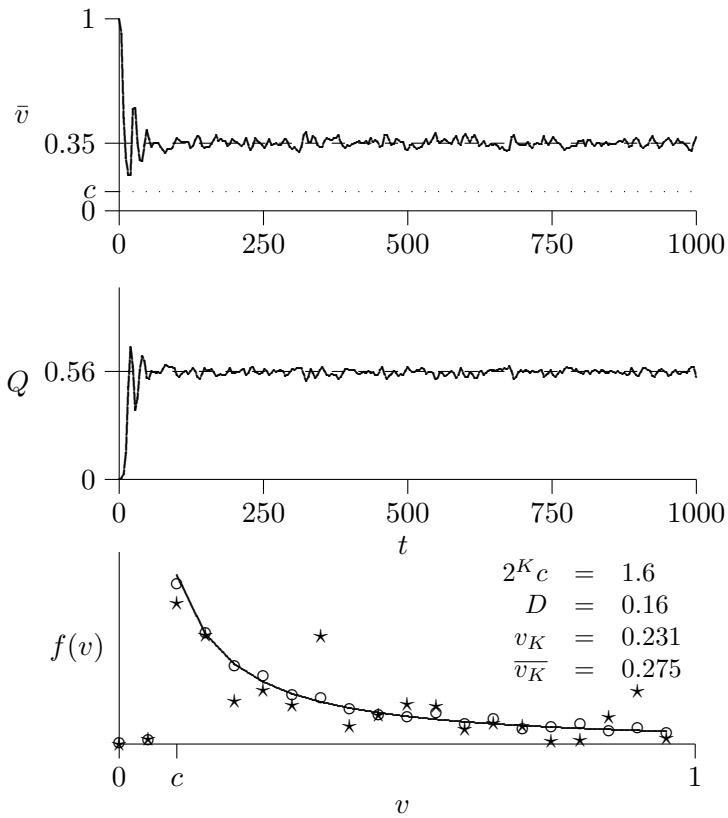

Figure 5: Simulation in which $c = 0.1$ and $K = 4$
Definitions are as in figure 4. In each generation there were 2000 groups of size 5. In the bottom panel, the aggregate distribution (circles) summarizes 4,201,045 individuals, and the final distribution (stars) summarizes 4653.

In the lower panels of these figures, the stars and circles represent empirical frequency distributions of the strategy variable v. The stars show distributions calculated from each simulation's final generation, while the circles show a distribution averaged over many generations. The solid line shows the predicted frequencies at the NE, as calculated from equation 10. The starred distribution fits the NE poorly in figure 4 but fairly well in figure 5. The long term average distribution (shown by circles) fits the NE well in both cases.

I did a series of such simulations, and in every case the system oscillated when $2^K c < 1$ but converged when $2^K c > 1$. These results are entirely consistent with condition 15, a fact that is somewhat surprising. Condition 15 does not guarantee stability in the simulations (even where it holds), because it only considers invasions that happen one at a time. When the population contains several invaders, the NE may be unstable even when inequality 15

holds. (For an example of this, see Maynard Smith [9, p. 185].) Yet no trace of this instability was discovered here. This may mean either that (1) the range of mixtures that can invade is so narrow that they hardly differ from the NE itself and are therefore not noticed, or (2) the mutational process generated such mixtures only rarely during the simulations.

5 Properties of a population playing the NE

In a world such as that of the simulation shown in figure 5, the NE is a good description of the population. How would such a population behave? Plugging equation 15 into equation 12 shows that the probability Q that a forager will fail to visit the resource must exceed 1/2. Thus, individual foragers will visit the resource less than half the time.

Surprising results emerge when one asks such questions as "How does the value of the harvested resource change with the number of competitors?" Intuition suggests that when the number of competitors is large, the resource will usually be harvested sooner and at a lower value. As we shall see, however, the model implies precisely the opposite.

Consider the probability distribution of the resource's value at the time it is harvested. If $K+1$ foragers are all playing strategy I, then the survival function of this new random variable is

$$s_K(v) = s(v)^{K+1} = (c/v)^{1+1/K} \tag{16}$$

This survival function gives the probability that the resource survives unharvested at least until its value is v. Equations 5 and 4 now give the hazard function and probability density:

$$h_K(v) = (1+1/K)/v \tag{17}$$

$$f_K(v) = \frac{(1+1/K)c^{1+1/K}}{v^{2+1/K}} \tag{18}$$

The probability that the resource is never harvested equals

$$Q_K := s_K(1) = c^{1+1/K}$$

For example, if $c = 1/2$ then the resource remains unharvested 1/4 of the time with two competitors but 1/2 the time with an infinite number. Apparently (but contrary to intuition), larger numbers of foragers leave more fruit on the tree.

The probability density of the value of resource when it is harvested, given that it is harvested at all, is $f_K(v)/(1-Q_K)$. The mean value of the resource when harvested is therefore

$$\bar{v}_K = (1-Q_K)^{-1}\int_c^1 v f_K(v)dv = \frac{(K+1)c(1-c^{1/K})}{1-c^{1+1/K}} \tag{19}$$

In the special cases of two competitors and of an infinite number, $\bar{v}_1 = 2c/(1+c)$ and $\bar{v}_\infty = -c(\ln c)/(1-c)$. With $c = 0.6$, these two cases give $\bar{v}_1 = 0.75$ and $\bar{v}_\infty = 0.77$. Note

Table 1: Theoretical and simulated values of v_K

K	v_K theory	v_K simulation	D
1	0.182	0.303	1.396
2	0.212	0.296	0.965
3	0.225	0.262	0.189
4	0.232	0.275	0.160
7	0.242	0.293	0.114
15	0.249	0.304	0.060

Note: $c = 0.1$. Each simulation involved 10,000 individuals and 1000 iterations.

that $\bar{v}_\infty > \bar{v}_1$. This means that on average (and contrary to intuition), the mean value of the resource when harvested increases with the number of competitors.

This result is also supported by simulations. Table 1 shows the results of six simulations representing an ascending series of values of K. In the table, D is a measure of the departure of the final frequency distribution from the Nash Equilibrium.[1] The simulations for $K = 1$ and $K = 2$ have large values of D, indicating that I provides a poor description. But D is small when $K \geq 3$, and in these simulations v_K increases with K. Thus, the simulations confirm that v_K increases with K.

To gain some intuition into this result, consider an alternative derivation. The mean benefit from a visit to the resource is $\bar{v}_K S$, where $\bar{v}_K$ is the mean value of the resource if harvested and S the probability that an individual who tries to harvest the resource will succeed. Equilibrium requires that

$$\bar{v}_K S = c \tag{20}$$

so that the expected gain from the berrypatch equals that from going fishing. (If this were not true, then it would pay to increase the fraction of one's time spent in the more productive activity, so we could not be at an equilibrium.) If c is constant, then any change in S must be offset by a change in $\bar{v}_K$. If S decreases with K, then $\bar{v}_K$ must increase. The more numerous the competitors, the more valuable the resource must be when harvested. An entirely separate argument thus leads to the same strange conclusion.

[1] Specifically, $D = \sum_{i=1}^{20} |p_{iI} - p_{iO}| + |Q_I - Q_O|$. Here, I have divided the interval from c to 1 into 20 equal subintervals and represent by p_{iI} the relative frequencies with which pure strategies inside interval i are played by strategy I. Similarly, p_{iO} denotes the relative frequency with which such strategies are played by individuals drawn from the population. Q_I and Q_O denote the relative frequencies with which the "go fishing" strategy is played by strategy I and by the population, respectively. D sums the absolute differences between the frequencies of strategy I and those in the simulated population.

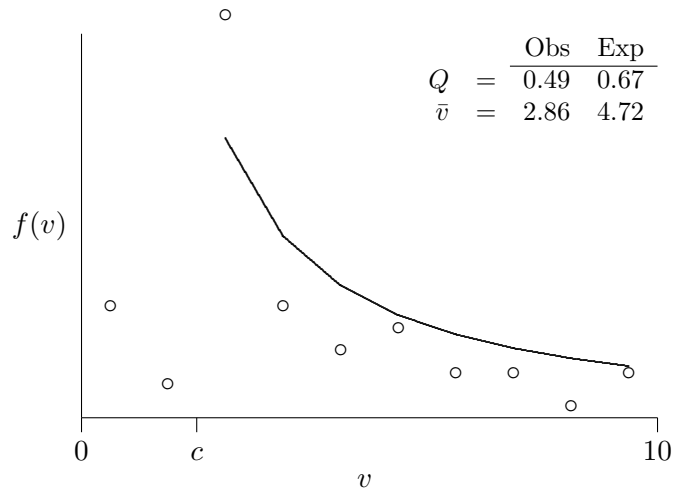

Figure 6: Experimental Results
Data from two experiments, each with 5 subjects and totaling 172 trials. Q is the frequency with which "go fishing" was played, and $\bar{v}$ is the mean strategy among the other plays.

6 Return to experimental data

Having reached an ambiguous conclusion from the theory, let us now return to the data that we started with. These data are re-plotted in figure 6 along with theoretical results for the case in which $K + 1 = 5$. The fit between observation and theory is not perfect: The observed value of Q is a little low (1/2 rather than 2/3), and the data give too much weight to values of v between 2 and 3. But on the other hand, the fit is not bad either. It is especially impressive in view of the fact that these data do not exclude the beginning of the game, when subjects were still learning.

7 Discussion

The classical theory of foraging ecology had little to say about game theory [11]. This made sense to the extent that the best strategy for one forager is independent of the strategies of others. But this independence seems implausible when several foragers compete for a single resource. The payoff from any given strategy surely depends on the strategies of others [6]. Thus, foraging ecology needs evolutionary game theory.

No pure strategy can be stable in this model, so any equilibrium must be mixed. Yet the only mixed Nash Equilibrium turns out to be unstable: there is always some mixed strategy that can invade. This result applies not only to the case in which the same mixed strategy is exhibited by each individual but also to the case in which the population contains a variety of strategies.

But this result is limited in an important way: It tells us only that the NE can be invaded by mixed strategies that are very similar to the NE itself. It need not be invadable by strategies that are substantially different. Thus, we have not excluded the possibility that the dynamics of this system remain close to the NE. If this were true, then the NE would be a good description of the population and its instability would be a mere technical curiosity. But while this possibility has not been excluded, it does not receive much support here either.

Some support is provided by the simulations even though the contestants there all played pure strategies. This is because a population of pure strategies that fails the test (inequality 14) for invasion by mixed strategies can be invaded by a mixture of pure-strategy mutants [9, p. 185]. This did not happen in the simulations. Consequently, appropriate mixtures of pure-strategy mutants must have arisen infrequently or not at all. This suggests that only a limited range of mixed strategies can invade the NE, in agreement with the conjecture of the previous paragraph.

It is also interesting that when one tallies the strategies played by human subjects in experiments, the result (figure 6) is reasonably close to the NE when inequality 15 is satisfied. This would make sense if my conjecture above were correct: In that case, genetic evolution should have shaped our psychological mechanisms so that we play something close to the NE. On the other hand, the experimental results also make sense on another hypothesis: Perhaps mixed strategies are hard for humans to employ. If so, then the experiments may resemble the simulations, where each individual played a pure strategy.

When behavior does approximate the NE, bizarre conclusions emerge. The more foragers there are, the less likely it is that a resource will be harvested at all and the higher the mean value of the resource when harvested. The model is unrealistically simple in assuming that the value of the resource is known with certainty and is the same for each competitor. Yet the simulations show that its predictions are remarkably accurate even when these simplifying assumptions are relaxed.

When the NE describes behavior, the resource will be ignored by at least half of the foragers who have access to it. Most harvesting attempts should occur when resources are barely ripe enough to offset costs and harvesting attempts will decline in frequency as the resource ripens. These predictions make sense of the ethnographic examples mentioned in the introduction, and also with the competition over strawberries in my back yard. It is no wonder that we all eat sour strawberries.

Acknowledgements

Much of this work was done during a sabbatical leave at the Research Centre of King's College, Cambridge, U.K. I am grateful for comments from T. Bergstrom, R. Boyd, E. Cashdan, E. Charnov, H. Gintis, A. Grafen, K. Hawkes, J. Hirshleifer, J.H. Jones, J. Kagel, Aaron McDonald, R. McElreath, Graham Mitchison, J. O'Connell, A. Roth, Paul

Smaldino, and Pontus Strimling. Finally, I thank those who helped by playing the berry-patch game.

A When does I resist invasion?

Each forager competes in a group with K others. Let P_0 denote the probability that a random forager competes with K individuals playing I and let P_1 denote the probability that he or she competes against one mutant playing strategy J and $K-1$ individuals playing I. I assume that mutants are rare so that the probability of competing against two or more mutants can be neglected; hence $P_0 + P_1 = 1$. The fitnesses of J-strategists and I-strategists are

$$\begin{aligned} W_J &= P_0\Pi(J, I^K) + P_1\Pi(J, J^1 I^{K-1}) \\ W_I &= P_0\Pi(I, I^K) + P_1\Pi(I, J^1 I^{K-1}) \end{aligned}$$

where the notation $\Pi(s, J^1 I^{K-1})$ refers to the expected payoff to strategy s when playing against K competitors, of whom 1 is playing J and the other $K-1$ are playing I. Strategy I is an ESS if and only if $W_I > W_J$. Since the definition of I implies that $\Pi(J, I^K) = \Pi(I, I^K) = 0$, W_I can exceed W_J only if

$$\Pi(I, J^1 I^{K-1}) > \Pi(J, J^1 I^{K-1}) \tag{21}$$

A.1 Invasion by pure strategies

Consider first the payoff to a mutant that plays pure strategy v in a group with one other mutant. The two mutants playing pure strategy v beat the $K-1$ I-strategists to the resource with probability $s(v)^{K-1}$, in which case they split the prize and each receive $v/2$. They pay a cost c whether they win or not. Thus,

$$\begin{aligned} \Pi(v, v^1 I^{K-1}) &= \frac{v}{2} s(v)^{K-1} \\ &= \frac{c}{2}\left(\frac{v}{c}\right)^{1/K} \\ \Pi(I, v^1 I^{K-1}) &= \int_c^v f(v) v s(v)^{K-1} dv + Qc \\ &= c\left(\int_c^v \frac{dv}{Kv} + Q\right) \\ &= c(\ln[(v/c)^{1/K}] + c^{1/K}) \end{aligned} \tag{22}$$

Stability requires that $\Pi(I, v^1 I^{K-1}) > \Pi(v, v^1 I^{K-1})$. This condition is equivalent to

$$\ln(v^{1/K}/Q) + Q - v^{1/K}/(2Q) > 0 \tag{23}$$

where Q is as defined in Eqn. 12. This inequality is satisfied if an only if $Q > 1/2$. This justifies inequality 15 for the case of mutants that play a pure strategy v such that $c \leq v \leq 1$.

The remaining case is that of a mutant that always fishes, never visiting the berry patch at all. Let us call this strategy F. The payoff to F is always c, no matter what its opponents do. Strategy I will resist invasion by F provided that

$$\Pi(I, F^1 I^{K-1}) > c$$

The payoff on the left can be found by setting $v = 1$ in Eqn. 22. The resulting expression is greater than c for all permissible values of c and K. Thus, strategy I always resists invasion by F. This completes the justification of inequality 15.

A.2 Invasion by a mixed strategy

This section will show that I is never stable against invasion by nearby mixed strategies. Subscripts will be used to distinguish quantities referring to different strategies: the survival and density functions of strategy I are denoted by s_I and f_I, and the corresponding functions of strategy J are s_J and f_J. The payoffs to I and J against groups with one J are

$$\Pi(I, J^1 I^{K-1}) = \int_c^1 f_I(v) v s_J(v) s_I(v)^{K-1} dv + Q_I c \tag{24}$$

$$\Pi(J, J^1 I^{K-1}) = \int_c^1 f_J(v) v s_J(v) s_I(v)^{K-1} dv + Q_J c \tag{25}$$

Adding and subtracting c to the right-hand side of both expressions, and making use of the fact that $v s_I(v)^{K-1} = c/s_I(v)$ leads to

$$\begin{aligned} \Pi(I, J^1 I^{K-1}) &= c\left[\int_c^1 f_I(v)\left(\frac{s_J(v) - s_I(v)}{s_I(v)}\right) dv + 1\right] \\ \Pi(J, J^1 I^{K-1}) &= c\left[\int_c^1 f_J(v)\left(\frac{s_J(v) - s_I(v)}{s_I(v)}\right) dv + 1\right] \end{aligned}$$

To measure the difference between the two payoffs, define

$$\begin{aligned} D &= [\Pi(I, J^1 I^{K-1}) - \Pi(J, J^1 I^{K-1})]/c \\ &= \int_c^1 [(s'_J - s'_I)(s_J - s_I)/s_I] dv \end{aligned} \tag{26}$$

where the arguments of s_J and f_J have been suppressed and f_I and f_J have been re-expressed as $-s'_I$ and $-s'_J$ (see equation 3). To prove that I is unstable, I must show that $D < 0$ for some J. Now $D = 0$ when $J = I$, for then D is then the difference between

two identical quantities. Consequently, I can prove that I is unstable by showing that D is greater when $J = I$ than otherwise. To this end, I use the calculus of variations to show that D reaches a local maximum where $s_J = s_I$.

The integrand within the definition of D can be written as

$$Z(s_J, s'_J, v) = (s'_J - s'_I)(s_J - s_I)/s_I$$

The function s_J that maximizes D must satisfy the Euler equation [4, p. 7],

$$\frac{\partial Z}{\partial s_J} - \frac{d}{dx}\left(\frac{\partial Z}{\partial s'_J}\right) = 0$$

Solving this differential equation with initial condition $s_J(c) = 1$ gives $s_J(v) = (c/v)^{1/K}$, which is identical to $s_I(v)$ (see equation 10). Thus, the only strategy that can possibly maximize D is strategy I, as defined in equations 9–11.

The calculus of variations requires functions with exogeneously determined endpoints. Consequently, I will stipulate that $s_J(1)$ and $s_I(1)$ are both equal to Q_I, as given in equation 12. If I can show that I cannot resist invasion by strategies that are constrained in this fashion, then it certainly cannot resist invasion by strategies chosen without constraint.

The Euler equation provides only a necessary condition and does not guaranteee that s_I maximizes D rather than minimizing it. To ensure that s_I is indeed a minimum, one must show that the "second variation" of D is positive. The second variation is [4, p. 35]

$$D_2 = \frac{\epsilon^2}{2!}\int_c^1 \left(t^2\frac{\partial^2 Z}{\partial s_J^2} + 2tt'\frac{\partial^2 Z}{\partial s_J \partial s'_J} + t'^2\frac{\partial^2 Z}{\partial {s'_J}^2}\right) dv$$

where ϵ is a small perturbation and t a function of v that is arbitrary except for the requirement that $t(c) = t(1) = 0$. In the present case, $\partial^2 Z/\partial s_J^2 = \partial^2 Z/\partial {s'_J}^2 = 0$, and $\partial^2 Z/(\partial s_J/\partial s'_J) = 1/s_I$. Thus, the integral in D_2 becomes

$$\int_c^1 \frac{tt'}{s_I} dv$$

Integrating by parts produces

$$\begin{aligned}\int_c^1 \frac{tt'}{s_I} dv &= \left.\frac{t(v)^2}{s_I(v)}\right|_c^1 - \int_c^1 t\left(\frac{t'}{s_I} - \frac{ts'_I}{s_I^2}\right) dv \\ &= 0 - \int \frac{tt'}{s_I} dv + \int_c^1 \frac{t^2 s'_I}{s_I^2} dv\end{aligned}$$

or

$$\int_c^1 \frac{tt'}{s_I} dv = \frac{1}{2}\int_c^1 \frac{t^2 s'_I}{s_I^2} dv \leq 0$$

The sign of the final term follows from the observations that $(t/s_I)^2 \geq 0$ and $s_I' \leq 0$ because survival functions cannot increase. This indicates that the second variation is negative, which implies that s_I maximizes D. This shows that I is not an ESS. It can be invaded by mixed strategies that are sufficiently similar to it.

Numerical example Consider a strategy J whose survival function is

$$\begin{aligned} s_J(v) &= s_I(v)(1 + bg(v)), \qquad \text{where} \\ g(v) &= (v - c)(1 - v)/(1 - c)^2 \end{aligned}$$

The function g is hump-shaped and implies that $s_J(1) = s_I(1) = Q_I$, as required by our previous assumptions. We must ensure that b is set to a value small enough that $s_J(v)$ has a negative derivative. Substituting into Eqns. 24 and 25 and setting $c = 1/2$, $K = 2$, and $b = 1/10$ gives $\Pi(I, J^1 I^{K-1}) = 0.5023689271$ and $\Pi(J, J^1 I^{K-1}) = 0.5023923175$. The second of these values is the larger, implying that J will invade a pure population of I. This confirms that I is *not* an ESS.

B When is a mixture of mixed types stable?

In the text, I have used the phrase "mixed ESS" to refer to

Interpretation 1 the mixed strategy played by each individual within a homogeneous population.

Here, I will argue that the results also apply when "mixed ESS" is interpreted as

Interpretation 2 the play of an individual chosen at random from a heterogeneous population containing several mixed strategies

provided that the payoff function satisfies a linearity condition that is described below.

Under interpretation 1, I is stable if and only if the following two conditions are met by all strategies J that differ from I:

$$\Pi(I, I^K) = \Pi(J, I^K) \tag{27}$$

$$\Pi(I, J^1 I^{K-1}) > \Pi(J, J^1 I^{K-1}) \tag{28}$$

My goal below is to show that these conditions are also necessary and sufficient to show that I is an evolutionary equilibrium (EE) under interpretation 2.

To simplify the exposition, I will ignore the distinction between probabilities and probability densities and will use the summation symbol $\sum$ to represent integration over continuous portions of the strategy space as well as summation over discrete portions.

Linearity condition Let A denote a possibly mixed strategy that plays pure strategy x with probability $p_A(x)$. The argument below will require that

$$\Pi(A, *) = \sum_x p_A(x)\Pi(x, *) \tag{29}$$

where the asterisk denotes a set of opponents whose types are unspecified. In words, the linearity condition requires that the payoff to A be a weighted average (or expectation) of payoffs to pure strategies that A may play.

I: A heterogeneous configuration of pure or mixed types. Consider a heterogeneous population containing n types of individual labeled $I_1, I_2, \ldots, I_n$, which have relative frequencies $w_1, w_2, \ldots, w_n$, and which play pure strategy x with probabilities $p_{I_1}(x), p_{I_2}(x), \ldots, p_{I_n}(x)$. I will use the symbol I to represent this configuration of pure or mixed strategies.

Payoffs involving I. An individual drawn at random from configuration I plays pure strategy x with probability

$$p_I(x) = \sum_i w_i p_{I_i}(x) \tag{30}$$

Payoffs to (or against) I will be intepreted as expected payoffs to (or against) random individuals chosen from configuration I. For example,

$$\Pi(I, B^K) = \sum_x p_I(x)\Pi(x, B^K) \tag{31}$$

and

$$\Pi(A, I^1 B^{K-1}) = \sum_x p_I(x)\Pi(A, x^1 B^{K-1}) \tag{32}$$

where A and B represent unspecified pure or mixed strategies. Given these definitions, the linearity condition 29 implies that

$$\Pi(I, B^K) = \sum_i w_i \Pi(I_i, B^K) \tag{33}$$

Thus, payoffs to configuration I may be expressed as a weighted average across the types that compose I.

If I is an EE then condition 27 must hold. To confirm this fact, suppose it to be false. If there is some pure strategy x for which $\Pi(x, I^K) > \Pi(I, I^K)$, then mutants playing that strategy would increase under natural selection. Therefore, if I is an EE then $\Pi(x, I^K) \leq \Pi(I, I^K)$. Now suppose that there is some pure x that is played by I with positive probability and for which $\Pi(x, I^K) < \Pi(I, I^K)$. Then since $\Pi(I, I^K)$ is the

average payoff to pure strategies (eqn. 31), and the payoff to x is below average, it follows that some other pure strategies must receive above-average payoffs. These types are being favored by selection, so once again I can't be an EE. Consequently, I cannot be an EE unless equation 27 is true.

If conditions 27 and 28 both hold, then any mutant invading I will decrease in frequency when rare. An invading mutant J will decrease in frequency if its fitness is less that the mean fitness within the population [3, p. 180]. This requires that

$$W_I > W_J \tag{34}$$

when J is rare, where W_J is the mutant's fitness and W_I is the mean fitness of individuals in configuration I. When mutant J has reached low frequency within the population, the fitnesses of the various types are

$$\begin{aligned} W_J &= P_0\Pi(J, I^K) + P_1\Pi(J, J^1 I^{K-1}) + \dots \\ W_{I_i} &= P_0\Pi(I_i, I^K) + P_1\Pi(I_i, J^1 I^{K-1}) + \dots \qquad (i = 1, 2, \dots, n) \end{aligned}$$

where P_i is the probability that a random individual (of any type) confronts i individuals of type J and $K - i$ individuals drawn at random from configuration I. The mean fitness of an individual in configuration I is

$$\begin{aligned} W_I &= \sum_i w_i W_{I_i} \\ &= P_0 \sum_i w_i \Pi(I_i, I^K) + P_1 \sum_i w_i \Pi(I_i, J^1 I^{K-1}) + \dots \\ &= P_0\Pi(I, I^K) + P_1\Pi(I, J^1 I^{K-1}) + \dots \end{aligned}$$

where the last line follows from equation 33. If J is sufficiently rare, we can ignore terms in $P_2, P_3, \dots$. And since condition 27 ensures that $\Pi(I, I^K) = \Pi(J, I^K)$, we can also drop the terms in P_0. Inequality 34 is then equivalent to inequality 28.

References

[1] Ian Moore Bennett. Barí Loricarid collection and the value of information: An application of optimal foraging theory. *Human Ecology*, 19:517–527, 1991.

[2] Nicholas G. Blurton Jones, Lars C. Smith, James F. O'Connell, Kristen Hawkes, and C.L. Kamuzora. Demography of the Hadza, an increasing and high density population of savanna foragers. *American Journal of Physical Anthropology*, 89:159–181, 1992.

[3] James F. Crow and Motoo Kimura. *An Introduction to Population Genetics Theory.* Harper and Row, New York, 1970.

[4] Charles Fox. *An Introduction to the Calculus of Variations*. Dover, New York, 1987.

[5] Herbert Gintis. *Game Theory Evolving*. Princeton University Press, Princeton, New Jersey, 2000.

[6] L.A. Giraldeau and T. Caraco. *Social Foraging Theory*. Princeton University Press, 2000.

[7] J. D. Kalbfleisch and R. L. Prentice. *The Statistical Analysis of Failure Time Data*. Wiley, New York, 1980.

[8] Dan Levin and James L. Smith. Equilibrium in auctions with entry. *The American Economic Review*, 84(3):585–599, June 1994.

[9] John Maynard Smith. *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, 1982.

[10] Martin J. Osborne and Ariel Rubinstein. *A Course in Game Theory*. MIT Press, Cambridge, Massachusetts, 1994.

[11] D. W. Stephens and J. R. Krebs. *Foraging Theory*. Princeton University Press, New Jersey, 1986.

[12] Peter Taylor and L. Jonker. Evolutionarily stable strategies and game dynamics. *Mathematicsl Biosciences*, 40:145–156, 1978.

[13] William Vickrey. Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37, 1961.

[14] Robert Wilson. A bidding model of perfect competition. *Review of Economic Studies*, 44:511–518, October 1977.

[15] E.C. Zeeman. Population dynamics from game theory. In *Proc. Int. Conf. Global Theory of Dynamical Systems*. Northwestern, Evanston, 1979.

[16] E.C. Zeeman. Dynamics of the evolution of animal conflicts. *Journal of Theoretical Biology*, 89:249–270, 1981.