



Do Evolutionary Processes Minimize Expected Losses?

DAVID B. FOGEL*[†] AND HANS-GEORG BEYER[‡]

**Natural Selection, Inc., 3333 N. Torrey Pines Ct., Ste. 200, La Jolla, CA 92037, U.S.A. and*

[‡]*University of Dortmund, Department of Computer Science, D-44221 Dortmund, Germany*

(Received on 9 February 2000, Accepted in revised form on 8 August 2000)

Evolution by variation and natural selection is often viewed as an optimization process that favors those organisms which are best adapted to their environment. This leaves open the issue of how to measure adaptation and what criterion is implied for optimization. This problem has been framed and analysed mathematically under the assumption that individuals compete to minimize expected losses across a series of decisions (e.g. choice of behavior), where each decision offers a stochastic payoff. But the fact that a particular analysis is tractable for a specified criterion does not imply the fidelity of that criterion. Computer simulations involving a version of the k -armed bandit problem can address the veracity of the hypothesis that individuals are selected to minimize expected losses. The results offered here do not support this hypothesis.

© 2000 Academic Press

Introduction

Evolution is commonly viewed as an optimizing process (Mayr, 1988; Alexander, 1996; and others). The question of what is being optimized has remained open to interpretation. One idea is that selection favors those behaviors which minimize expected losses over a series of decisions, where each decision receives a stochastic payoff (Holland, 1975). Such a criterion is convenient mathematically because it is amenable to an analysis of optimal strategies. But such mathematics cannot address the principal question of the suitability of the criterion itself; it merely assumes the criterion as given. However, the aptness of this criterion can be modeled by computer simulations in which individuals compete for sur-

vival based on the reward they receive while employing different strategies for sampling from a k -armed bandit.

The k -armed bandit serves as a familiar analogy that can provide insight into animal behavior in diverse environments (Alexander, 1996). Variations of the analogy are typical in optimality models that presume individuals can adopt alternative behavioral strategies, each having a range of payoffs with certain likelihoods, much like pulling a one-armed bandit (a slot machine). These payoffs can be measured in terms of food obtained, shelter, reproductive success, or other suitable standards (Krebs *et al.*, 1978; Green, 1980; Bateson & Kacelnik, 1997). In traditional evolutionary game theory, expected payoffs translate linearly into reproductive success or “fitness” (Maynard Smith, 1982). When interest is focused on the specific behaviors that provide the foundation for that reproductive success, however, the relationship between behaviors

[†]Author to whom correspondence should be addressed.
E-mail: dfogel@natural-selection.com; www.natural-selection.com

and fitness may be nonlinear. Alternative behavioral strategies can be compared in light of a chosen mathematical criterion, with those that are found to be optimal in turn being compared to those observed in nature. It is hoped that with appropriate abstraction, this mathematical device can provide insight into the fundamental dynamics that underlie the observed behaviors.

One widespread abstraction of the k -armed bandit problem to adaptation in natural settings is found in the canonical genetic algorithm (Holland, 1975), in which the problem of adaptation is framed as a series of decisions on how to best allocate trials to alternative bandits in light of payoffs received from previous trials. The principal assumption underlying the mathematical framework of genetic algorithms is that the natural selection minimizes expected losses while sampling from alternative bandits [described as “schemata” (Holland, 1975)]. Under this assumption, Holland (1975) offered what was presumed to be an optimal sampling strategy of devoting an exponentially increasing number of trials to the observed best bandit (the one with the greatest observed average payoff). Recent independent analyses (Rudolph, 1997; Macready & Wolpert, 1998) have, however, proved that the development by Holland is mathematically flawed: the preferred strategy of the genetic algorithm is not, in truth, optimal for the criterion of minimizing expected losses.

Rather than seek to fix this circumstance by discovering alternative strategies that are indeed optimal for this criterion, a more fundamental question can be raised: does selection in fact favor those behavioral strategies which seek to minimize expected losses, and if not, what are the conditions that determine which tactics will be favored by selection? A first answer to this question can be obtained by studying a two-armed bandit problem using a simulated population of competing strategies that are subject to random variation and selection. The hypothesis that selection favors minimizing expected losses can then be tested statistically by examining the strategies that survive over a large number of generations. As will be seen, the results of these simulations do not offer general support for this hypothesis.

Methods

A population of N individuals was constructed where each faced the choice of sampling from either of two slot machines. The first machine offered a payoff that was Gaussian distributed with a mean of 1.0 and standard deviation (S.D.) of 1.0, while the second was also Gaussian distributed but with zero mean and a standard deviation that was parameterized by the symbol σ_2 . After each individual sampled from either of the bandits, all individuals were ranked in order of decreasing payoff and the subset K of these with the greatest payoffs were selected to generate subsequent progeny. In this way, K represented the carrying capacity of the environment. This process was then iterated over several thousands of “generations”.

The behavior of each individual was defined by a single parameter, p_i , $i = 1, \dots, N$, which corresponded to the probability that it would sample from the first bandit (i.e. the one with the greater mean). This protocol is typical of phenotypic optimality models where the underlying genetics of a particular behavior is abstracted out of consideration (Maynard Smith, 1982). A surplus of offspring was generated from surviving individuals through slight random variation of each parent’s parameter. Each surviving individual was given an equal probability of being selected to generate each next offspring ($i = 1, \dots, N$). Specifically, each offspring’s parameter, p'_i was set equal to its parent’s parameter p_i with the addition of zero mean Gaussian noise with standard deviation of 0.01. This choice was deemed reasonable for representing a small amount of persistent random variation. If any offspring’s p'_i became greater than 1.0 or less than 0.0 it was set to the limit it exceeded, thereby maintaining its interpretation as a probability. Note that each parent could generate more than one offspring.

Consideration was given to cases where the maximum population size, N , was significantly larger than the carrying capacity, K . This represents the case of stringent selection pressure where a great percentage of offspring do not survive to reproduce and the reproductive strategy of generating a tremendous surplus of offspring is adopted, rather than having parental investment be constrained to only a small number of

offspring. Experiments were conducted for the cases where (1) $K = 1$ and $N = 100$, and (2) $K = 100$ and $N = 10\,000$ (i.e. on average, each parent generated 100 offspring) at various settings of σ_2 ranging from 0.01 to 5.0 at selected intervals. Each individual was defined to have a maximum lifespan of only one generation, that is, all parents were removed from the population at each iteration, and only the best K out of the N offspring were selected for further reproductive attention. In both the (1, 100) and (100, 10 000) cases, the initial population of N individuals was selected with each individual's probability parameter being chosen uniformly at random over the interval $[0, 1]$.

Attention was focused on the mean of all surviving parents' probability parameters at each generation. To avoid initial transient effects, data were recorded from generations 100 000 to 10 million for the case of (1, 100) and from generations 20 000 to 120 000 for the case of (100, 10 000). The mean probability parameters over these generations were then averaged to generate a single datum representing the mean probability of choosing the first bandit for each investigated setting of σ_2 . Under the hypothesis that selection favors strategies that minimize expected losses, the anticipated results would indicate a strong tendency to sample from the first bandit regardless of σ_2 because it has the higher average payoff (Holland, 1975).

Results

Figure 1(a) and (b) show the results for both cases. The figures indicate a clear shift in optimal (i.e. selected) behavior away from sampling the bandit with the higher mean as the standard deviation of the second bandit was increased beyond a particular threshold in the range of $1 < \sigma_2 < 2$. (The threshold point is analysed mathematically below.) The natural outcome of these simple evolutionary systems was a selection for risky behavior even when the average payoff for that risk was lower than that offered by the less variable option. Note also that even at very low values of σ_2 there was a saturation of the mean probability of selecting the first bandit at values significantly lower than 1.0. That is, even when the second bandit had very low variance,

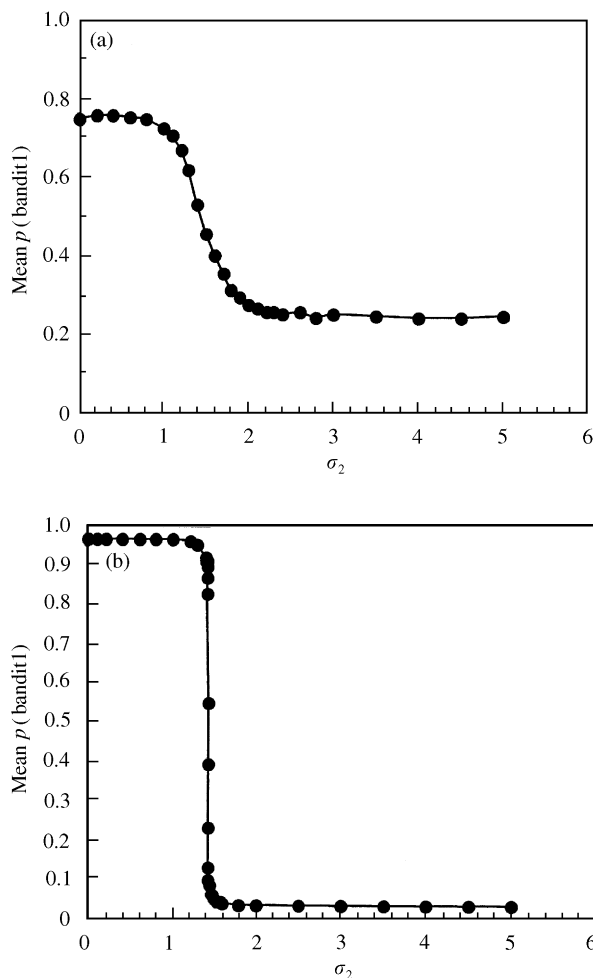


FIG. 1. The mean probability of choosing the bandit with greater average payoff as a function of the standard deviation of the bandit with lower average payoff. (a) $K = 1$ and $N = 100$, and (b) $K = 100$ and $N = 10\,000$. Sampling was conducted from generation 100 000 to 10 million for case (a), and 20 000 to 120 000 for case (b), so as to avoid transient effects that might result from the uniform initialization of strategies. Under the hypothesis that selection would favor strategies that minimize expected losses, the anticipated result would be a flat line indicating complete certainty of choosing the first bandit regardless of the value of σ_2 . In contrast, the results indicate not only a lack of convergence to strategies which always choose the first bandit when σ_2 is small but also a consistent shift that favors strategies which choose the second bandit with lower mean payoff as σ_2 is increased. Each datum is the average on the order of 10^5 or 10^7 trials and thus the 95% confidence limits around each point are not visible at the scale shown. The results indicate statistically significant evidence rejecting the hypothesis that selection favors strategies which minimize the expected losses ($p \ll 10^{-6}$).

there was insufficient selection pressure to drive the mean probability of choosing the bandit with a higher average payoff to complete certainty. Even under these conditions, selection did not

favor strategies that minimized expected losses exclusively.

Analysis

It is of interest to determine a mathematical description of the relationship between the mean probability of choosing the first bandit given the means and standard deviations of the two bandits ($\mu_1, \mu_2, \sigma_1, \sigma_2$, respectively), and the values of N and K . Closed-form mathematical analysis was undertaken but the result was intractable. Analysis of the case of N offspring where the carrying capacity is $K = 1$ also appears intractable except for the trivial case of $N = 2$. It is, however, possible to offer useful formulas for identifying the transition point where the mean probability of choosing the first bandit shifts from high to low values as a function of the parameters of the two bandits.

It was noted by observing simulations for various values of N and K that the variability of the (selected) p_i parameters (recall, p_i is the parameter of the i -th individual defining their probability of choosing the bandit with the greater mean) decreases with the increasing N and K . That is, for N large and $1 \ll K \ll N/2$, the distribution of p_i becomes so concentrated that a useful first approximation is to replace the random variate defined by the p_i events by its expected value, say τ . The expected numbers of individuals sampling from the two bandits are then τN and $(1 - \tau)N$, respectively. Assuming the bandits are described by the random variables

$$Y \sim \mu_1 + \sigma_1 N(0, 1), \quad Z \sim \mu_2 + \sigma_2 N(0, 1), \quad (1)$$

respectively. The τN samples from Y and the $(1 - \tau)N$ samples from Z can be ordered according to the standard notation from order statistics:

$$f_{1:\lambda} \leq f_{2:\lambda} \leq \dots \leq f_{\lambda-1:\lambda} \leq f_{\lambda:\lambda}, \quad (2)$$

where λ is the number of samples to be ordered with respect to the payoff values f that they earned. The values received from each arm can be ordered as

$$Y_{1:\tau N} \leq Y_{2:\tau N} \leq \dots \leq Y_{\tau N:\tau N},$$

$$Z_{1:(1-\tau)N} \leq Z_{2:(1-\tau)N} \leq \dots \leq Z_{(1-\tau)N:(1-\tau)N}. \quad (3)$$

The K best individuals (i.e. those with the largest payoffs) are selected. Let τ' be the portion of these individuals that come from Y and $(1 - \tau')K$ be the resulting number of selected individuals that come from Z .

Assuming without loss of generality that a sample from Z is the worst of the K selected individuals. Then,

$$Y_{\tau N - \tau'K + 1:\tau N} \geq Z_{(1-\tau)N - (1-\tau')K + 1:(1-\tau)N}$$

and

$$Y_{\tau N - \tau'K:\tau N} < Z_{(1-\tau)N - (1-\tau')K + 1:(1-\tau)N}. \quad (4)$$

Under stationary conditions, the portion τ' of individuals selected from Y must be equal to the portion τ of those actually sampled from Y . Thus, setting $\tau' = \tau$ yields:

$$Y_{\tau(N-K) + 1:\tau N} \geq Z_{(1-\tau)(N-K) + 1:(1-\tau)N}$$

and

$$Y_{\tau(N-K):\tau N} < Z_{(1-\tau)(N-K) + 1:(1-\tau)N}. \quad (5)$$

While these inequalities can only hold statistically, it is reasonable to assume that they hold approximately for the expected values of the sampled random variables. That is, in the limit case as N gets large, stationarity is characterized by

$$E[Y_{\tau(N-K) + 1:\tau N}] = E[Z_{(1-\tau)(N-K) + 1:(1-\tau)N}]. \quad (6)$$

With Y and Z defined as in eqn (1):

$$\mu_1 + \sigma_1 c_{\tau(N-K) + 1:\tau N} = \mu_2 + \sigma_2 c_{(1-\tau)(N-K) + 1:(1-\tau)N}, \quad (7)$$

where $c_{a:b}$ is the expected value of the $a:b$ order statistics of the $N(0, 1)$ normal variate (Arnold *et al.*, 1992):

$$c_{a:b} = a \binom{b}{a} (2\pi)^{-0.5} \int_{-\infty}^{\infty} x \exp(-x^2/2) (\Phi(x))^{a-1} \times (1 - \Phi(x))^{b-a} dx, \quad (8)$$

where $\Phi(x)$ is the cumulative distribution function of the $N(0, 1)$ standard normal random variable ($c_{a:b}$ is the expected value of the a -th highest number drawn from a sample of b). Given the bandit parameters μ_1, μ_2, σ_1 , and σ_2 in eqn (7), σ_2 can be expressed as a function of τ :

$$\sigma_2 = \frac{\mu_1 - \mu_2 + \sigma_1 c_{\tau(N-K)+1:\tau N}}{c_{(1-\tau)(N-K)+1:(1-\tau)N}}. \quad (9)$$

Equation (9) can be further simplified by considering the asymptotic behavior of $c_{a:b}$. From the theory of order statistics (Arnold *et al.*, 1992), for a and b sufficiently large:

$$c_{a:b} \simeq \Phi^{-1}\left(\frac{a}{b+1}\right), \quad (10)$$

where $\Phi^{-1}(x)$ is the inverse of the cumulative distribution function of $N(0, 1)$. Thus, eqn (9) becomes

$$\sigma_2 \simeq \frac{\mu_1 - \mu_2 + \sigma_1 \Phi^{-1}[\tau(N-K) + 1/\tau N + 1]}{\Phi^{-1}[(1-\tau)(N-K) + 1/(1-\tau)N + 1]}. \quad (11)$$

For $N \rightarrow \infty$:

$$\sigma_2 \simeq \sigma_1 + \frac{\mu_1 - \mu_2}{\Phi^{-1}[(N-K)/N]}. \quad (12)$$

This surprisingly simple formula can predict the phase transition point well, especially as N gets large and K being not too small. For example, for the case shown in Fig. 1(b) with $N = 10000$ and $K = 100$, eqn (12) yields $\sigma' = 1.4292$ as the threshold value of σ_2 . We have tested eqns (9), (11), and (12) in simulations using $N = 1000$ and $K = 5, 10, 50, 100, 200, 300, 400$, and 450 (note that $1 \ll K \ll N/2$ as was assumed in the above derivation) and found good agreement between the predicted and observed behavior [as is also observed in Fig. 1(a)].

In addition to predicting the threshold of transitioning between the two bandits, analysis can be performed to indicate the expectation of the average payoff to the surviving parents as a function of σ_2 for large N and $1 \ll K \ll N/2$. As seen in Fig. 2 for the case of $N = 1000$ and




FIG. 2. The observed average payoff $\langle f \rangle$ (which estimates $E[\langle f \rangle]$) earned by the surviving parents for the case of $N = 1000$ and $K = 10$ as a function of σ_2 found by simulation. The data were obtained by averaging $\langle f \rangle$ over 200 000 generations after an initial phase of 10 000 generations. The progress coefficient needed in eqns (13) and (17) is $c_{10/10,1000} \approx 2.64695$. As detailed in the text, analysis explains the observations in terms of two distinct regimes. The first occurs when σ_2 is less than the threshold value where sampling switches between the first and second bandit. In this regime, the surviving parents have nearly constant average payoffs regardless of σ_2 . When σ_2 increases beyond the threshold value, the average payoff earned by the surviving parents increases linearly with σ_2 . The two lines superimposed on the data indicate the expected mean payoff in each regime derived by analysis. The lines are extended to provide convenient comparison.

$K = 10$, the empirical mean parental fitness takes on two distinct regimes: when $\sigma_2 < \sigma'$, the mean threshold value when the observed behavior shifts, the mean parental fitness is constant but when $\sigma_2 > \sigma'$, the mean parental fitness increases linearly. For the assumed conditions on N and K , when $\sigma_2 < \sigma'$, $\tau \approx 1$. Assuming that almost all surviving parents sample from $Y \sim \mu_1 + \sigma_1 N(0, 1)$, their expected average payoff is

$$\begin{aligned} E[\langle f \rangle] &= E\left[\frac{1}{K} \sum_{i=1}^K (\mu_1 + \sigma_1 N(0, 1))_{N-i+1:N}\right] \\ &= \mu_1 + \sigma_1 \frac{1}{K} \sum_{i=1}^K c_{N-i+1:N} \\ &= \mu_1 + \sigma_1 c_{K/K,N}, \end{aligned} \quad (13)$$

where

$$c_{K/K,N} = \frac{1}{K} \sum_{i=1}^K c_{N-i+1:N}, \quad (14)$$

which is known in the literature of evolutionary algorithms as the so-called *progress coefficient*. This can be expressed by the integral (Beyer, 1995)

$$c_{K/K,N} = \frac{N-K}{2\pi} \binom{N}{K} \int_{-\infty}^{\infty} e^{-t^2} (\Phi(t))^{N-K-1} \times (1 - \Phi(t))^{K-1} dt, \quad (15)$$

which has asymptotic behavior:

$$c_{K/K,N} \simeq \frac{N}{K} \frac{1}{\sqrt{2\pi}} \exp \left[\frac{1}{2} \left(\Phi^{-1} \left(\frac{N-K}{N} \right) \right)^2 \right]. \quad (16)$$

Note that eqn (13) does not vary as a function of σ_2 (i.e. it exhibits constant behavior).

For $\sigma_2 > \sigma'$, $\tau \approx 0$. Assuming almost all samples of the surviving parents are from $Z \sim \mu_2 + \sigma_2 N(0, 1)$ yields

$$\begin{aligned} E[\langle f \rangle] &= E \left[\frac{1}{K} \sum_{i=1}^K (\mu_2 + \sigma_2 N(0, 1))_{N-i+1:N} \right] \\ &= \mu_2 + \sigma_2 \frac{1}{K} \sum_{i=1}^K c_{N-i+1:N} \\ &= \mu_2 + \sigma_2 c_{K/K,N}. \end{aligned} \quad (17)$$

This confirms that the expected average parental fitness increases linearly with σ_2 for $\sigma_2 > \sigma'$.

Deriving formulae for the mean probability of selecting from the first bandit in the limiting cases of $\sigma_2 \rightarrow 0$ and ∞ as a function of N and K remains for future work.

Discussion

Several design choices were made for the current simulations that may affect the observed shift in behavior toward the more risky strategy of sampling from the bandit with a lower mean but larger variance. Chief among these is the stringency of selection pressure coupled with a large surplus of offspring. This situation necessitates risky behavior because as the possible reward for taking that risk becomes greater, the conservative strategy of opting for the less variable but greater average payoff becomes unten-

able. Given a sufficient number of risk takers, a sufficient subset of those will get lucky and gain a payoff that is larger than is likely when choosing conservatively. Most risk takers will be losers, but at the same time most of the winners in this lottery will also be risk takers.

It is of interest to identify conditions in natural settings that are similar to those incorporated in the models studied here. The situation of a high selection pressure and large surplus of offspring is not uncommon in nature (*r*-selection). Individual mortality rates are often difficult to estimate, but there have been some assessments. The daily rate of larval mortality in northern anchovies has been estimated between 16 and 20% (Hewitt, 1981; Blaxter & Hunter, 1982). It has also been estimated that 70% of the eggs of Atlantic herring in a patch off the Canadian Coast were eaten by predatory flounders (one examined flounder contained 16 000 eggs) (Tibbo *et al.*, 1963). These rates concern only immature fishes and therefore the percentage of individuals that survive to maturity and further go on to reproduce must be even lower. Similar observations can be made for a great number of insect species and flora [e.g. ferns generate billions of spores annually (Gould & Gould, 1989)], in contrast with, say, mammals which are often *K*-selected; however, even in this latter case, predation is often a serious threat (Cheney *et al.*, 1988; Boesch, 1991; Stanford *et al.*, 1994), and sexual selection may be very stringent, with few of the males actually taking part in reproduction (e.g. with the intense competition between male elephant seals less than one-third of the beach-resident males copulate during a breeding season and most matings are accomplished by only a few males (McCann, 1981; Andersson, 1994); moreover, only about 10% of male pups are still alive when they are mature enough to compete seriously for rookeries with other males (Gould & Gould, 1989).

The idea that the variability of payoffs plays a role in choosing optimal behavioral strategies has been examined in depth by animal behaviorists in different settings (Stephens & Krebs, 1986; McNamara & Houston, 1992). For example, Caraco *et al.* (1980) offered juncos choices between being fed with a fixed or variable number of seeds, with both choices constrained to have the same mean number of seeds (i.e. 2 vs.

0 or 4). The juncos were observed to prefer the more variable choice after being deprived of food for a prolonged period of time. Similar tendencies toward more variable payoffs when under stress, even though the mean payoff for each choice was identical, have been observed in other studies (Caraco *et al.*, 1990; Cartar & Dill, 1990). This is *prima facie* evidence contradicting the explanation that the behaviors have evolved to minimize expected losses, for if this criterion held true there would be no expected bias for making either choice given identical mean payoffs.

The results presented here goes further by indicating that risky behavior can be preferred even when the mean payoff is sacrificed if the potential rewards are sufficiently high and sufficiently likely. Moreover, selection may not completely eliminate behaviors that yield a lower expected reward even when the variance of this subpar payoff is small. Quantifying the relationships between the difference in mean reward between two or more alternative choices and their associated variability remains for future work, but the results are clear: evolution does not, in general, minimize the expected losses. When payoffs do not translate linearly into reproductive success, minimizing expected losses does not imply maximizing reproductive success. Conclusions derived from models of adaptation in natural and artificial systems that adopt this criterion should be viewed with appropriate skepticism.

We thank L. Altenberg, R. W. Anderson, P. J. Angeline, W. Atmar, K. Chellapilla, G. B. Fogel, L. J. Fogel, W. G. Macready, and V. W. Porto for technical assistance and review of the manuscript. We also thank the referees for their valuable comments. The second author acknowledges support by the DFG as a Heisenberg Fellow under grant Be 1578/4-1.

REFERENCES

- ALEXANDER, R. M. (1996). *Optima for Animals*, revised edition. Princeton, NJ: Princeton University Press.
- ANDERSSON, M. (1994). *Sexual Selection*. Princeton, NJ: Princeton University Press.
- ARNOLD, B. C., BALAKRISHNAN, N. & NAGARAJA, H. N. (1992). *A First Course in Order Statistics*. New York: John Wiley.
- BATESON, M. & KACELNIK, A. (1997). Starlings' preferences for predictable and unpredictable delays to food. *Anim. Behav.* **53**, 1129–1142.
- BEYER, H.-G. (1995). Toward a theory of evolution strategies: on the benefits of sex—the $(\mu \setminus \mu, \lambda)$ -theory. *Evol. Comp.* **3**, 81–111.
- BLAXTER, J. H. S. & HUNTER, J. R. (1982). The biology of clupeoid fishes. In *Advances in Marine Biology*, Vol. 20, (Blaxter, J. H. S., Russell, J. & Young, C., eds), pp. 1–203. New York: Academic Press.
- BOESCH, C. (1991). The effects of leopard predation on grouping patterns in chimpanzees. *Behaviour* **117**, 220–242.
- CARACO, T., MARTINDALE, S. & WHITTAM, T. S. (1980). An empirical demonstration of risk-sensitive foraging preferences. *Anim. Behav.* **28**, 820–830.
- CARACO, T., BLANCKENHORN, W. U., GREGORY, G. M., NEWMAN, J. A., RECER, G. M. & ZWICKER, S. M. (1990). Risk-sensitivity: ambient temperature affects foraging choice. *Anim. Behav.* **39**, 338–345.
- CARTAR, R. V. & DILL, L. M. (1990). Why are bumble bees risk-sensitive foragers? *Behav. Ecol. Sociobiol.* **26**, 121–127.
- CHENEY, D. L., SEYFARTH, R. M., ANDELMAN, S. J. & ANDELMAN, L. P. (1988). Reproductive success in vervet monkeys. In: *Reproductive Success* (Clutton-Brock, T. H., ed.), pp. 382–402. Chicago: Chicago University Press.
- GOULD, J. L. & GOULD, C. G. (1989). *Sexual Selection*. New York: Scientific American Library.
- GREEN, R. F. (1980). Bayesian birds: a simple example of Oaten's stochastic model of optimal foraging. *Theor. Popul. Biol.* **18**, 244–256.
- HEWITT, R. (1981). The value of pattern in the distribution of young fish. In: *The Early Life History of Fish* (Lasker, R. & Sherman, K., eds), Vol. 178, pp. 229–236. *Rapports et Proces-Verbaux des Reunions. Conseil International pour l'Exploration de la Mer*.
- HOLLAND, J. H. (1975). *Adaptation in Natural and Artificial Systems*. Ann Arbor, MI: University of Michigan Press.
- KREBS, J. R., KACELNIK, A. & TAYLOR, J. P. (1978). Test of optimal sampling by foraging great tits. *Nature* **275**, 27–31.
- MACREADY, W. G. & WOLPERT, D. H. (1998). Bandit problems and the exploration/exploitation tradeoff. *IEEE Transactions on Evolutionary Computation* **2**, 2–22.
- MAYNARD SMITH, J. (1982). *Evolution and the Theory of Games*. Cambridge, U.K.: Cambridge University Press.
- MAYR, E. (1988). *Toward a New Philosophy of Biology: Observations of an Evolutionist*. Harvard, MA: Belknap.
- MCCANN, T. S. (1981). Aggression and sexual activity of male Southern elephant seals, *Mirounga leonina*. *J. Zool.* **195**, 295–310.
- MCMANARA, J. M. & HOUSTON, A. I. (1992). Risk-sensitive foraging: a review of the theory. *Bull. Math. Biol.* **54**, 355–378.
- RUDOLPH, G. (1997). Reflections on bandit problems and selection methods in uncertain environments. In *Proc. Seventh International Conference on Genetic Algorithms* (Bäck, T., ed.), pp. 166–173. San Mateo, CA: Morgan Kaufmann.
- STANFORD, C. B., WALLIS, J., MATAMA, H. & GOODALL, J. (1994). Patterns of predation by chimpanzees on colobus monkeys in Gombe National Park, 1982–1991. *Am. J. Phys. Anthropol.* **94**, 213–228.
- STEPHENS, D. W. & KREBS, J. R. (1986). *Foraging Theory*. Princeton, NJ: Princeton University Press.
- TIBBO, S. N., SCARRATT, D. J. & MCMULLION, P. W. G. (1963). An investigation of herring (*Clupea harengus* L.) spawning using free-diving techniques. *J. Fisheries Res. Board Canada* **20**, 1067–1079.

AUTHOR QUERY FORM

HARCOURT PUBLISHERS

JOURNAL TITLE: JTBI
ARTICLE NO. : 20002166

DATE: 30/08/2000

Queries and / or remarks

Manuscript Page/line	Details required	Author's response
10/13	Carter or Cartar?	