frontiers in
COMPUTATIONAL NEUROSCIENCE

How should prey animals respond to uncertain threats?

Joel Zylberberg[1,2]* and Michael Robert DeWeese[1,2,3]

[1] Department of Physics, University of California, Berkeley, CA, USA
[2] Redwood Center for Theoretical Neuroscience, University of California, Berkeley, CA, USA
[3] Helen Wills Neuroscience Institute, University of California, Berkeley, CA, USA

A prey animal surveying its environment must decide whether there is a dangerous predator present or not. If there is, it may flee. Flight has an associated cost, so the animal should not flee if there is no danger. However, the prey animal cannot know the state of its environment with certainty, and is thus bound to make some errors. We formulate a probabilistic automaton model of a prey animal's life and use it to compute the optimal escape decision strategy, subject to the animal's uncertainty. The uncertainty is a major factor in determining the decision strategy: only in the presence of uncertainty do economic factors (like mating opportunities lost due to flight) influence the decision. We performed computer simulations and found that *in silico* populations of animals subject to predation evolve to display the strategies predicted by our model, confirming our choice of objective function for our analytic calculations. To the best of our knowledge, this is the first theoretical study of escape decisions to incorporate the effects of uncertainty, and to demonstrate the correctness of the objective function used in the model.

Keywords: decision making, agent-based modeling, uncertainty, evolution, escape decision, probabilistic automata, constrained optimization

1 INTRODUCTION

Prey animals frequently assess their surroundings to identify potential threats to their safety. If an animal does not flee soon enough in the presence of a predator (type I error), it may be injured or killed. If it flees when there is no legitimate threat (type II error), it wastes metabolic energy, and loses mating or foraging opportunities (Nelson et al., 2004; Creswell, 2008). However, predators can be camouflaged and prey animals have limited sensory systems, so escape decisions must often be made with imperfect information.

Previous studies have not investigated how escape decisions might be affected by prey animals' degree of certainty about their environment. Indeed, the predominant assumption in the field appears to be that this uncertainty is not important, and that so long as the prey animal knows the most likely state of the environment (or the expected value of the state), they can still make economically optimal decisions. We question this assumption.

We explicitly consider the animal's uncertainty in our model and subsequently demonstrate that, when a prey animal knows the environmental state with certainty, the optimal decision strategy is simply to flee whenever a threat is present. This strategy is independent of any "economic" factors – predator lethality, predator frequency, loss of mating opportunities, etc. When the state of the environment is less certain, the animal is bound to make errors, and the optimal balance between type I and II errors is determined by economic factors. This is in contrast with previous theoretical studies (Ydenberg and Dill, 1986; Broom and Ruxton, 2005; Cooper, 2006; Cooper and Frederick, 2007, 2010; and others), which have assumed that prey animals have perfect knowledge of their surroundings (or, equivalently, that the uncertainty is unimportant, as discussed above) and that their decisions are made on purely economic grounds. Our result suggests that uncertainty may play a key role in making economics relevant in decision making.

It has been observed that prey animals increase their flight initiation distance (predator–prey distance at which they flee) when intruders begin their approach from farther away (Blumstein, 2003). This observation challenges strictly economic decision models (although Cooper et al., 2009 has modeled this by assuming the prey animal has multiple "risk" functions and chooses one based on the predator behavior, which includes starting distance) but may have a simple interpretation: when intruders approach from farther away, the animal has longer to detect and assess the threat, thus fleeing sooner. This "increased information" interpretation is also consistent with the observations that prey animals increase their flight initiation distance when an intruder has a higher approach speed (Cooper, 2006), and that odors (Apfelbach et al., 2005; Ylonen et al., 2007) and shapes (Hemmi and Merkle, 2009) are important in eliciting defensive behaviors – the approach speed may indicate that the intruder is likely to attack, while the odor and morphology may indicate that the intruder is a potentially dangerous one.

The current models of escape decisions (Ydenberg and Dill, 1986; Broom and Ruxton, 2005; Cooper, 2006; Cooper and Frederick, 2007, 2010; and others) have yet to incorporate factors such as olfaction and intruder morphology into their models. Such factors are difficult to include in economic models. For example, what is the cost associated with a particular smell?

Inspired by these observations, we propose a new approach for studying escape decisions in prey animals, namely that they are engaged in a decision-theoretic process, wherein they must decide, with imperfect information, whether the current environment is likely to pose a significant enough threat to their safety that they should flee. This view is supported by observations of active risk assessment behaviors in prey animals (Schaik et al., 1983; Creswell et al., 2009; Hemmi and Pfeil, 2010).

Sih (1992) has studied the decision to re-emerge from a burrow after flight, when the animal does not know if the predator is still present. Sih's work is the closest in spirit to the current study, but it does not address how the initial escape decision is affected by uncertainty. No previous study satisfactorily addresses the issue of determining what objective function, when maximized, accurately predicts the escape decision strategy that is selected by evolution.

We demonstrate through computer simulation that animals subject to predation naturally evolve to display the strategy predicted by our model, confirming our choice of objective function.

2 METHODS AND MODELS

2.1 ANALYTICAL CALCULATIONS

As a starting point, we will assume that the prey animal chooses the strategy S that maximizes its genetic contribution to subsequent generations (see Janetos and Cole, 1981; Parker and Smith, 1990 for criticism and discussion of optimality models), defined by the intrinsic rate of increase in prevalence of strategy S in the population (Hairston et al., 1970; Parker and Smith, 1990); $r(S) = [N(S)^{-1}]$ $dN(S)/dt$, where $N(S)$ is the number of animals adopting strategy S, and t is time (note that we use the symbol r for the same quantity that Hairston et al., 1970 call m). Since the large r strategies grow more quickly in terms of number of adherents, the population should evolve toward the strategy that maximizes r. We later verify this assumption. We stress that our objective function is *rate of reproduction* and not "survival of the fittest." A genotype that leads to long-lived animals, who fail to reproduce, is unlikely to significantly increase in prevalence over time.

We formulate a probabilistic automaton (Rabin, 1963) model of the life of a prey animal (**Figure 1**) to compute $r(S)$. The model operates in discrete time, thus we are using the approximation $r = [N(S)^{-1}]dN(S)/dt \approx [N(S)^{-1}]\Delta N(S)/\Delta t$. In every time step, every

animal follows a complete path through the graph, beginning in the starting state ("start"), and ending either back at the start, or in death.

In each time step, the animal assesses a potential threat. For concreteness, we imagine the animal asking "Is that object likely to try to kill me?" Animals that do not flee from a real threat may be killed by a predator, while those that do flee, escape. Those animals that are not killed by predators may mate, and they may or may not die of causes other than predation.

We group potential threats into discrete "zones" in predator–prey distance; **Figure 1** illustrates an example with four such zones. The model can utilize continuous distances without affecting our conclusions. The object the animal assesses may or may not be a real threat – the animal does not have access to the ground truth. We explicitly incorporate this uncertainty by assigning probability p_i of correct threat detection in zone i, leading to flight, and probability q_i of a false positive in zone i, leading to flight; q_i and p_i are related by the receiver–operator-characteristic (ROC) curve [$p_i = f_i(q_i)$; see **Figure 2**].

Qualitatively, our automaton model captures many features pertinent to real prey animals. Effects like periodicity of mating opportunities and threat frequency, maturation periods, learning during the lifetime of the animal (Hemmi and Merkle, 2009; Rao, 2010), and sexual reproduction (as opposed to asexual), are omitted in the interest of simplicity, but our automaton could be amended to incorporate these considerations. We have confirmed with a computer simulation that our results are unchanged when the animals undergo sexual, rather than asexual reproduction (results not shown).

Figure 2 presents an example ROC curve, for the case in which the animal makes its choice based on a single, scalar parameter (the "score"). The problem of choosing an escape strategy amounts to
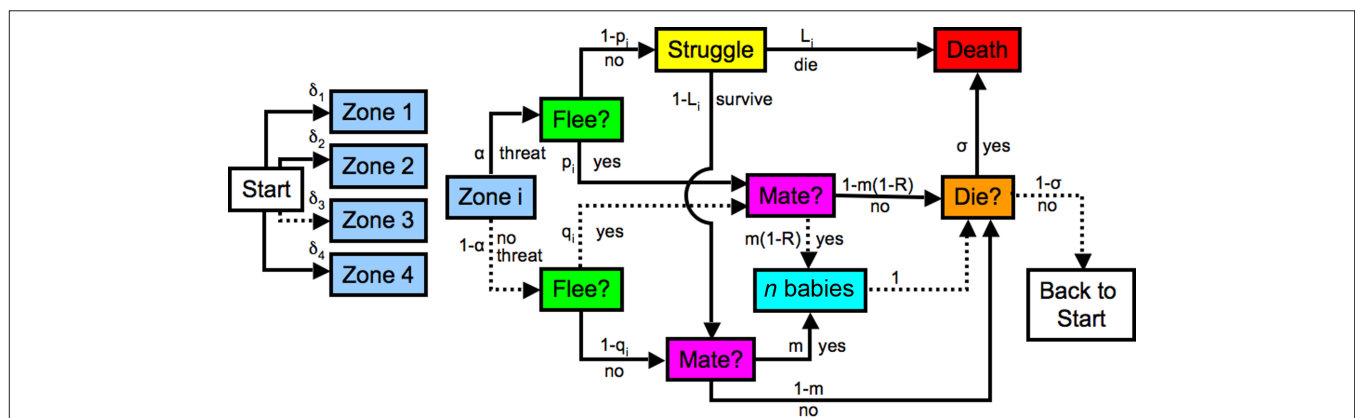

FIGURE 1 | Probabilistic automaton model of the life cycle of a prey animal. At each time step, every animal begins in the state "start," and follows a complete path, ending either back at the start, or in death. Each arrow is labeled with the conditional probability that the given event occurs (die, survive, etc.), once the animal reaches the box at the tail of that arrow. The animal spots a potential threat in zone i with probability δ_i. Four zones (groupings by predator–prey distance) are shown in the diagram. The threat is real with probability α. If there is a threat the prey animal flees with probability p_i. Those animals that do not flee are killed with probability L_i, while those that do flee always escape. The animals that neither flee nor die mate with probability m, producing n progeny.

The animals that do flee suffer a reduced mating rate of $m(1 − R)$. In order to keep the population stable, some randomly selected animals are killed at the end of the time step with probability σ. The probability of any path is obtained by multiplying the conditional probabilities of each subsequent step. A sample path is illustrated by the dashed arrows in the diagram: The animal spots a potential threat in zone 3. This is not a real threat, but, with its imperfect information, the animal incorrectly decides to flee. It then mates, producing n progeny, which are added to the population for the next time step. The animal does not succumb to disease, starvation, or other non-predation-related causes of death, and lives on to the next time step. The probability of this path is $\delta_3(1 − \alpha)q_i m(1 − R)(1 − \sigma)$.

choosing where on the ROC curve the decision rule should lie. It can be specified in each zone either by p_i, q_i, or a threshold τ_i (see **Figure 2**).

To simplify our notation, we first define the variables a,b,c to be the probabilities, in a given time step, of being killed by a predator, of not fleeing, and not being killed by a predator, and of fleeing, respectively. By tracing paths in **Figure 1**, we find that

$$
\begin{aligned}
a &= \sum_i \delta_i \alpha [1 - p_i] L_i \\
b &= \sum_i \delta_i \left(\alpha [1 - p_i][1 - L_i] + [1 - \alpha][1 - q_i] \right) \\
c &= \sum_i \delta_i \left(\alpha p_i + [1 - \alpha] q_i \right) \\
1 &= a + b + c \\
1 &= \sum_i \delta_i
\end{aligned}
\tag{1}
$$

where α is the probability that a given object is actually a threat, and L_i is the probability that failure to immediately flee a predator, initially in distance zone i, will be lethal.

The expectation value of the objective function ($r = [N(S)^{-1}] dN(S)/dt \approx [N(S)^{-1}]\Delta N(S)/\Delta t$) is given by multiplying the result of a given outcome by the probability of that outcome, and summing over all possible outcomes. These outcomes are as follows: animals that do not flee will mate with probability m, producing n progeny, while those that do flee will suffer a reduced mating rate $m(1 - R)$, producing n progeny. Animals that die are removed from the population, thus the value of this outcome

FIGURE 2 | Connection between ROC curves (p vs. q) and probability distributions. The probability distributions of some "score" parameter, conditioned on the presence (solid curve) or absence (dashed curve) of a threat are shown. A possible interpretation for this score is that it is the output of a neural network that assesses all of the information available to the animal, in an attempt to infer the danger of a given object. One possible decision threshold, τ, is indicated, whereby the animal decides to flee from objects with scores above the threshold, and not to flee from those with scores below the threshold. By varying the threshold, the animal can alter the correct detection probability p, indicated by the area under the "threat" distribution (solid curve) to the right of the threshold; at the same time, varying the threshold also affects the false positive probability q, given by the area under the "no threat" distribution (dashed curve) to the right of the threshold. The same threshold determines both p and q; they are related by the ROC curve $p = f(q)$ (inset).

is -1. Those animals that are not already killed by the predator die with probability σ, which again has a value -1. Thus, our expectation value is

$$
\begin{aligned}
E\left[r(\{q_i\})\right] &\approx b(\{q_i\})mn + c(\{q_i\})mn[1 - R] - a(\{q_i\}) - \sigma(1 - a\{q_i\}) \\
&\approx \left[1 - a(\{q_i\})\right]mn - c(\{q_i\})Rmn - a(\{q_i\}) - \sigma.
\end{aligned}
\tag{2}
$$

The second line follows from the first since $a + b + c = 1$, and, since both σ and a are expected to be small, the product σa can safely be ignored.

Since the animals in our model assess one potential threat per unit time, the size of the "time steps" in our model is fairly short (seconds, or possibly minutes). In the real-world, we expect that actual threats are relatively uncommon (for example, the probability of encountering a real threat in any given short time period is small): α should be a relatively small quantity. Thus, a is small for real prey animals. Furthermore, since the time steps are fairly short, the probability σ of dying from starvation or disease in any time step is quite small. Thus, our $\sigma a << 1$ approximation (above) is reasonable.

The anticipated escape response threshold maximizes the expectation value of the objective function $E[r(\{q_i\})]$, subject to the constraints $p_i = f_i(q_i)$ imposed by the ROC curves for each zone.

In the standard fashion (Boas, 2006), we utilize the method of Lagrange multipliers by defining a Lagrange function $\Omega = E[r] + \Sigma_i \xi_i(p_i - f_i(q_i))$. The set $\{\xi_i\}$, then, is then our set of (unknown) Lagrange multipliers, and we optimize by solving (for all i)

$$
\begin{aligned}
\frac{\partial \Omega}{\partial p_i} &= \frac{\partial E\left[r(\{q_i\})\right]}{\partial p_i} + \xi_i = 0 \\
\frac{\partial \Omega}{\partial q_i} &= \frac{\partial E\left[r(\{q_i\})\right]}{\partial q_i} - \xi f_i'(q_i) = 0 \\
\frac{\partial \Omega}{\partial \xi_i} &= p_i - f_i(q_i) = 0.
\end{aligned}
\tag{3}
$$

The last of these equations enforces the constraint. The first two equations yield

$$
\begin{aligned}
\xi_i &= -\alpha L_i[1 + mn] + \alpha Rmn \\
\xi_i f_i'(q_i) &= -Rmn[1 - \alpha].
\end{aligned}
\tag{4}
$$

And the solution to our optimization problem is (for $\xi_i \neq 0$)

$$
\begin{aligned}
f_i'(q_i) &= \frac{Rmn[1 - \alpha]}{\alpha\left(mn[L_i - R] + L_i\right)} \\
p_i &= f_i(q_i).
\end{aligned}
\tag{5}
$$

Note that, were all of the threats in one distance zone, Eq. 5 still yields the optimal result. Thus, the globally optimal solution consists of making the locally optimal decision for each zone, as one might expect.

For an explicit computation of where the decision threshold should lie, we require information about the ROC curve. As an example, we assume that the score is distributed as $g_i(z|\text{danger}) = \mathcal{N}(0,1)$ in the presence of danger in zone i, where $\mathcal{N}(\mu, \Sigma)$ represents a
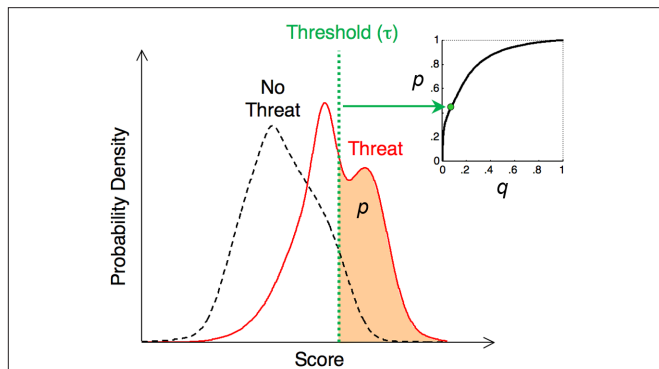
Gaussian (or normal) distribution with mean μ and SD Σ. Now let the distribution of scores in the absence of danger in zone i be $g_i(z|\text{no danger}) = \mathcal{N}(\omega_i,1)$ for some $\omega_i < 0$.

Note that, given a Gaussian-distributed variable y with arbitrary mean μ and variance Σ, we can choose to operate on the variable $x = (y - \mu)/\Sigma$, which will be distributed as $\mathcal{N}(0,1)$. Thus, within the realm of Gaussian-distributed scores, we are losing no generality by considering $g_i(z|\text{danger}) = \mathcal{N}(0,1)$. Choosing the variance of "no threat" score distribution to be the same as that of the distribution conditioned on the presence of danger does entail a loss of generality, but it simplifies the analysis greatly and thus we do it for the purposes of this example. Given the distributions, we can define the values (p_i, q_i) as a function of the decision threshold τ_i. Let the animal decide that it is in danger for $z > \tau_i$, and that it is not for $z \le \tau_i$. Then

$$p_i = f_i\left(q_i(\tau_i)\right) = \int_{\tau_i}^{\infty} \frac{dz}{\sqrt{2\pi}} e^{-z^2/2}$$

$$q_i(\tau_i) = \int_{\tau_i}^{\infty} \frac{dz}{\sqrt{2\pi}} e^{-(z-\omega_i)^2/2}. \tag{6}$$

We need the derivative $f_i'(q_i)$ to implement the results of our optimization calculation. Using the chain rule,

$$f_i'\left(q_i(\tau_i)\right) = \frac{dp_i}{d\tau_i}\frac{d\tau_i}{dq_i}$$

$$= \exp\left(\frac{\omega_i^2 - 2\omega_i\tau_i}{2}\right). \tag{7}$$

Therefore, the optimal threshold for the i^{th} zone is

$$\tau_i = \frac{\omega_i}{2} - \frac{1}{\omega_i}\ln\left(\frac{Rmn(1-\alpha)}{\alpha\left[mn(L_i - R) + L_i\right]}\right). \tag{8}$$

It is clear that as $|\omega_i|$ increases (more obvious threats and thus less uncertainty), the second term, which contains all of the economic factors about the environment, becomes less important in determining the decision threshold.

The decreasing importance of the "economic" term with increasing $|\omega|$ is not true for all possible score distributions. We have demonstrated that this conclusion does apply to Gaussian distributions. Indeed, it also applies to any unimodal distribution in the exponential family $g(z) \propto e^{-\lambda|z|^\nu}$ with even $\nu \ge 2$.

This can be seen by noting that, if the distribution of scores in the presence of a threat is $g(z|\text{threat}) \propto e^{-\lambda|z|^\nu}$, and in the absence of a threat is $g(z|\text{no threat}) \propto e^{-\lambda|z-\omega|^\nu}$, then the derivative of the ROC curve $f'(q)$ is given by

$$f'(q(\tau)) = \exp\left(\lambda|\tau|^\nu - \lambda|\tau - \omega|^\nu\right)$$

$$\rightarrow \lambda^{-1}\ln\left[f'(q)\right] = |\tau|^\nu - |\tau - \omega|^\nu. \tag{9}$$

Solving for τ is hard for general ν. Consider, for example the case where $\tau \ge 0$ and $\tau \ge \omega$. Then we see that $\lambda^{-1}\ln[f'(q)] = \tau^\nu - (\tau - \omega)^\nu$. For $\nu = 1$, this yields no solution for τ because the derivative of the ROC curve is independent of τ. This is peculiar to the exponential distribution, which is a pathological case in this sense.

For even ν, with no restrictions on τ, we see that $\lambda^{-1}\ln[f'(q)] = \tau^\nu - (\tau - \omega)^\nu$ (the absolute value signs disappear for even ν). Expanding $(\tau - \omega)^\nu$ using binomial theorem, we find

that $\lambda^{-1}\ln[f'(q)] = \tau^\nu - \tau^\nu - \sum_{j=1}^{\nu}\binom{\nu}{j}\tau^{\nu-j}(-\omega)^j$. Now, the τ^ν terms cancel, and we can divide through by one power of ω, yielding $\lambda^{-1}\omega^{-1}\ln[f'(q)] = \sum_{j=1}^{\nu}\binom{\nu}{j}\tau^{\nu-j}(-\omega)^{j-1}$. As in the Gaussian ($\nu = 2$) case, we see that increasing $|\omega|$ de-weights the economic $f'(q)$ term.

Thus, we can be assured that the decreasing importance of the economic term with increasing $|\omega|$ is true for all unimodal exponential distributions of the form $g(z) \propto e^{-\lambda|z|^\nu}$ with even $\nu \ge 2$. For $\nu > 1$ and values of ν that are not even integers, there are some regimes in which the leading-order terms in τ still cancel, however it is difficult to prove that our result holds in the most general case.

We note that, while it simplified our automaton model and our notation, nowhere was it necessary to assume that the danger occurs in discrete zones in distance. One could instead utilize a continuous distance measure by considering an infinitely large set of possible values of i, with each one corresponding to a particular point in space.

2.2 SIMULATION EXPERIMENTS

To verify that our objective function is the one selected for by natural evolution, we perform a computer simulation of a population of prey animals subject to predation.

Our simulation contains a population of animals whose life cycles are described by the probabilistic automaton model (**Figure 1**). At each time step of the simulation, the animals are considered one-by-one. A pseudo-random number generator determines whether a prey animal will see a real threat (with probability α) or not. The threats are all in the same distance zone, since this simplifies the simulation, and we have shown that the optimal solution for many zones is to use the locally optimal solution in each separate zone (Eq. 5).

The animal is then presented with a "score" variable, with which it makes its decision. As in our analytic example, the scores are randomly drawn from the $\mathcal{N}(0,1)$ distribution if the threat is real, or from the $\mathcal{N}(\omega,1)$ distribution if the threat is fake. If the "score" is above the animal's threshold, it chooses to flee. Otherwise it does not. The determination of which animals get to mate, or get killed by a predator is also done with a pseudo-random number generator, and follows the description in **Figure 1**.

Those animals that do mate produce n progeny. Each offspring has a decision threshold that is equal to its parent's, plus Gaussian noise of mean zero and fixed (small) SD. This variation allows the population to explore the strategy space. The population in our simulation thus has the two key features (heritability, and variability) that allow for evolution.

At the end of every time step, the population is trimmed so that it does not get too large. This is done by killing random individuals, thus inducing no selection pressure. This is represented by the value σ in our automaton model.

We initialize the simulation with a population of animals whose decision thresholds are drawn from a uniform distribution.

3 RESULTS AND DISCUSSION
3.1 THE OPTIMAL DECISION STRATEGY DEPENDS ON THE ENVIRONMENT AND VARIES WITH THE ANIMAL'S UNCERTAINTY ABOUT THE STATE OF THE ENVIRONMENT

The strategy that maximizes the expectation value of r, subject to the $p_i = f_i(q_i)$ constraints imposed by the ROC curves, is given by (for all zones i)

$$f_i'(q_i) = \frac{Rmn(1-\alpha)}{\alpha\left[(L_i - R)mn + L_i\right]}. \tag{10}$$

Figure 2 shows that low thresholds yield small derivatives $f_i'(q_i)$ (for low thresholds, a small increase in τ increases q much more than p), and *vice versa*. We observe that increased predator density and lethality leads to more timid prey animals (low τ), while increased reproductive flight cost Rmn leads to bolder ones.

To investigate the influence of uncertainty on the decision strategy, we consider a specific example: the decision is made based on a single "score" variable z, which is distributed as $g_i(z|\text{danger}) = \mathcal{N}(0,1)$ in the presence of danger in zone i, where $\mathcal{N}(a,b)$ is a Gaussian (or normal) distribution with mean a and SD b.

This score may be, for example, the output of a neural network that takes into account all of the information available to the animal, including information about the predator behavior, odor, morphology, etc. The use of a single score for the decision can be understood as a dimensionality reduction step: the high-dimensional sensory data is reduced to a single scalar value, upon which the decision can be based. In the case of an animal with a "command neuron" (e.g., the Mauthner cell; Rock et al., 1981; Roberts, 1992; Zottoli and Faber, 2000; Korn and Faber, 2005), the score we refer to may be related to the membrane potential, which is a function of all the synaptic inputs to that cell from the sensory processing network.

Let the distribution of scores in the absence of danger in zone i be $g_i(z|\text{no danger}) = \mathcal{N}(\omega_i, 1)$ for some $\omega_i < 0$. The absolute value of ω_i defines the reliability of the information available to the animal: larger $|\omega_i|$ implies more reliable information. Let the animal decide that it is in danger for $z > \tau_i$, and that it is not in danger for $z \leq \tau_i$. The optimal threshold τ_i for zone i is

$$\tau_i = \frac{\omega_i}{2} - \frac{1}{\omega_i}\ln\left[f_i'(q_i)\right], \tag{11}$$

and $f_i'(q_i)$, which incorporates all of the economic factors in the probabilistic automaton model, is given by Eq. 5.

For large $|\omega_i|$, the dependence on the logarithmic term is small and, for $f_i'(q_i) > 0$ (which is true, for example, when $L_i > R$ and $\alpha \neq 1$), the threshold is $\tau_i \approx \omega_i/2$. This strategy resembles a maximum likelihood estimator (MLE), which, in this example, would be given by $\tau_i = \omega_i/2$. As $|\omega_i|$ decreases, the economic factors become more important in determining the threshold τ_i. This conclusion (demonstrated in **Figure 3**) is independent of the details of our probabilistic automaton model. The specifics of the automaton determine the optimal $f_i'(q_i)$, but Eq. 11 shows us that, independent of $f_i'(q_i)$, the strategy still changes from maximum likelihood to economic cost reduction, as the amount of uncertainty in the information increases.

This result is true for Gaussian-distributed score variables, but is not true for all distributions. However, it is straightforward to prove that the result holds for all unimodal distributions in the exponential family $g(z) \propto e^{-\lambda|z|^\nu}$ for even $\nu \geq 2$. According to the central limit theorem, most variables that are weighted averages of many random components are Gaussian-distributed. Thus, our conclusion is likely to be applicable to many real-world examples.


FIGURE 3 | The importance of economic factors in decision making increases with rising uncertainty about the environment. The departure of the optimal decision threshold (τ) from a MLE (described by $\tau = \omega/2$) is shown as a function of ω, the displacement between the means of the score distributions for threats and non-threats. The result is shown for several different values of $f'(q)$, which contains all the economic factors, and quantifies how bold (large values) or timid (small values) the strategy is (see text). For large $|\omega|$, threats are easily identified, and the strategies all converge to a maximum likelihood decision strategy: flee if and only if danger is more likely than not. As the uncertainty increases (small $|\omega|$), the strategies diverge in a manner dictated by the economic factors.

3.2 AN *IN SILICO* POPULATION OF PREY ANIMALS EVOLVES TO DISPLAY THE DECISION STRATEGY PREDICTED BY OUR ANALYTICAL CALCULATIONS

To verify our choice of objective function, and the approximations made in our calculation, we performed a computer simulation of a population of prey animals subject to predation. Unlike previous work (Floreano and Keller, 2010), we did not define an objective function in our simulation: the animals in our simulation had no indication of what we thought they should be accomplishing with their escape strategy. They simply mated, died of predation, and were killed by non-predation-related causes. Those animals that did mate produced children whose escape thresholds were copies of the parent's threshold, with added Gaussian noise. We investigated how the evolutionarily favored escape response threshold varied as a function of the parameters of their life cycle, and various properties of their predators.

Some of the results of this experiment are shown in **Figure 4**. For the results in the scatter plot, the simulation was repeated many times. For each run of the simulation, the parameters ($\alpha, L, m, n, R, \omega$) were randomly selected from i.i.d. uniform distributions with the ranges specified in **Table 1**. The population average threshold was recorded after each time step, and the result shown on the plot is the average over the last 10^3 time steps. This reduces the variance of the results that stems, in part, from the variance of the children with respect to their parents, and in part from the relatively small population ($N = 2 \times 10^3$) that was used in the simulation. This variance is depicted in **Figure 4**.

The results of the simulation (**Figure 4**) demonstrate that a population of animals whose life cycle is well-described by the automaton model in **Figure 1** will naturally evolve to display the strategy defined in Eqs. 10 and 11. Much of the challenge in applying our method to real prey animals will be in accurately modeling their life cycle with a probabilistic automaton model.
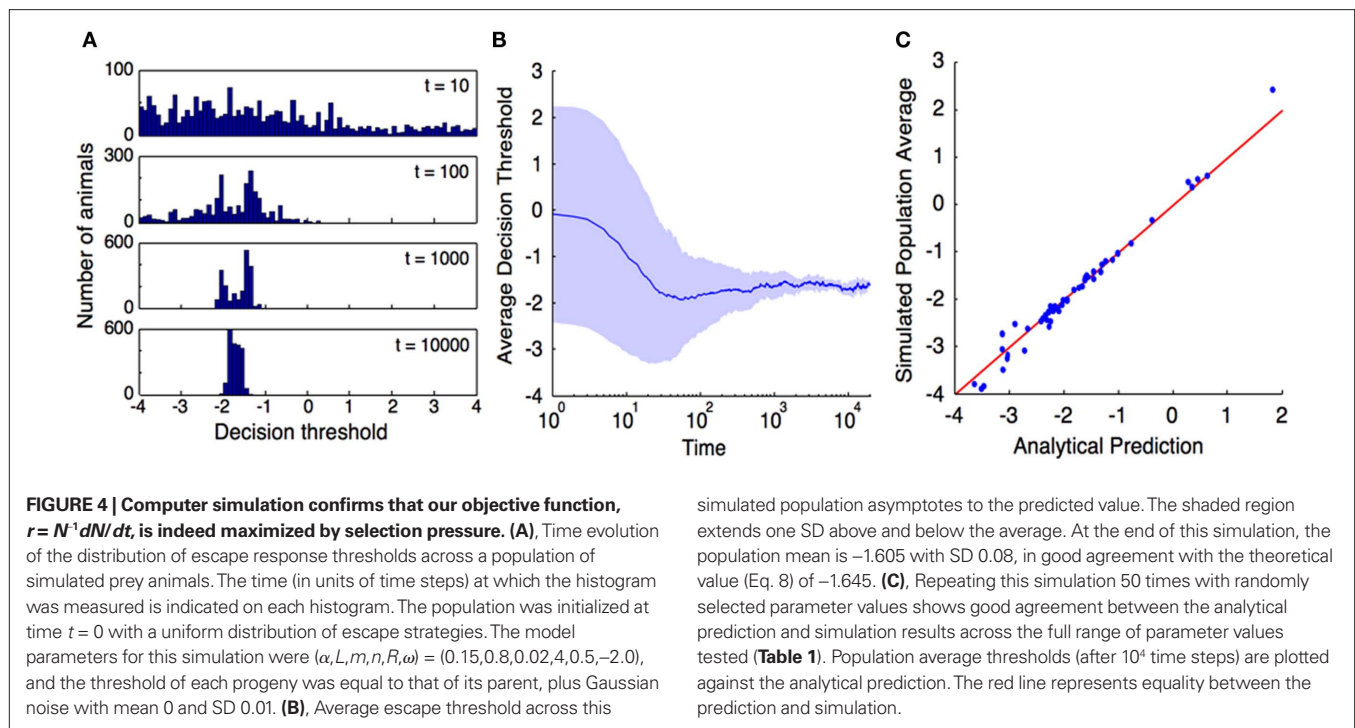
FIGURE 4 | Computer simulation confirms that our objective function, $r = N^{-1}dN/dt$, is indeed maximized by selection pressure. (A), Time evolution of the distribution of escape response thresholds across a population of simulated prey animals. The time (in units of time steps) at which the histogram was measured is indicated on each histogram. The population was initialized at time $t = 0$ with a uniform distribution of escape strategies. The model parameters for this simulation were $(\alpha, L, m, n, R, \omega) = (0.15, 0.8, 0.02, 4, 0.5, -2.0)$, and the threshold of each progeny was equal to that of its parent, plus Gaussian noise with mean 0 and SD 0.01. **(B)**, Average escape threshold across this

simulated population asymptotes to the predicted value. The shaded region extends one SD above and below the average. At the end of this simulation, the population mean is −1.605 with SD 0.08, in good agreement with the theoretical value (Eq. 8) of −1.645. **(C)**, Repeating this simulation 50 times with randomly selected parameter values shows good agreement between the analytical prediction and simulation results across the full range of parameter values tested (**Table 1**). Population average thresholds (after 10^4 time steps) are plotted against the analytical prediction. The red line represents equality between the prediction and simulation.

Table 1 | Ranges of variables used in the simulation.

Variable	Range
α	[0.05, 0.6]
L	[0.5, 1.0]
m	[0.01, 0.5]
n	[3, 8]
R	[0.1, 0.7]
ω	[−7.0, −0.5]

We stress that we made a specific choice of objective function for our analytic calculation, but that objective function was not available to the animals in our simulation. Had we made a different choice of objective function, our analytic calculations would have yielded different results, and those would necessarily not have been in agreement with the simulation results.

For example, choosing longevity as an objective function, one would choose the strategy that maximizes lifetime. Given the structure of our automaton model, that strategy is clearly to flee all of the time; $\tau = -\infty$, regardless of the model parameters. That result is clearly in disagreement with our simulation results. Thus, we argue that the results of our simulation support our chosen objective function.

It has previously been conjectured (Cooper and Frederick, 2007) that, because the correct objective function is unknown, and prey animals have uncertain information about the environment, quantitative behavioral predictions are impossible. We have addressed both of these issues: the correct objective function, while hard to compute for real prey animals (much information is required to correctly estimate r), is known, and we have explicitly incorporated the effects of imperfect information into our decision model.

We conclude that, given sufficiently accurate and detailed (probabilistic) information about the life cycle of an animal, it may be possible (although difficult) to make quantitative behavioral predictions.

DISCUSSION AND CONCLUSION

We have found that a prey animal's uncertainty about threats in its environment has a profound effect on the optimal escape strategy. Moreover, computer simulations of the evolution of populations of animals subject to predation demonstrate that the objective function we assumed for our analytic calculations is, indeed, optimized by selection pressure.

Interesting work has modeled the learning process in the presence of uncertainty, in the context of optimal decision making (Rao, 2010). Our results focus instead on instinctual responses, and do not explicitly incorporate learning over the lifetime of the animal. Clearly, in the real-world, both innate and learned behaviors are important. We leave the issue of combining these two response types for future work.

Whereas much previous work (Ydenberg and Dill, 1986; Blumstein, 2003; Broom and Ruxton, 2005; Cooper, 2006; Cooper and Frederick, 2007, 2010) has involved determining flight initiation distances, our model does not do so explicitly. In our model, the animal simply flees when the possibility of danger exceeds some threshold, the value of which is determined by the level of uncertainty, and, when that uncertainty is not small, by economic factors. However, we do have the ability to infer how such a strategy might vary when assessing threats at different distances.

We expect that nearby threats will be more conspicuous: $|\omega|$ should be a decreasing function of distance. Thus, the economic factors are more important for potential threats at large distances compared to small. Consequently, those economic factors that

make the strategy more timid (lower threshold) will increase the flight initiation distance – they make the optimal threshold lower at large distances, but do not affect the small distance threshold as strongly.

Similarly, when an intruder initiates its approach from further afield, the prey animal has more time to gain information about it. Thus, at a greater distance, the animal can correctly assess the threat, leading to a larger flight initiation distance, as observed in real prey animals (Blumstein, 2003).

Finally, when an odor or shape is presented to the animal that is associated with common predators, the score of the intruder will be far from the mean of the "no threat" distribution. Thus, defensive behavior is likely to be trigered.

Our decision-theoretic model for prey escape strategy can thus account for several observed behaviors (Blumstein, 2003; Apfelbach et al., 2005; Cooper, 2006) in a natural way. Indeed, to the best of our knowledge, ours is the first model to account (qualitatively) for the influence of all of these factors on escape decisions.

We propose that approaches based on optimal performance in the face of imperfect information are likely to be useful for studying further aspects of escape decisions in prey animals, as they have been in other areas of biology such as mate selection (Benton and Evans, 1998; Luttbeg and Warner, 1999), house-hunting (Marshall et al., 2006), cellular-level decision processes (Perkins and Swain, 2009), and chemotaxis (Adler and Wung-Wai, 1974; Bialek and Setayeshgar, 2005). Our approach is easily generalized to include other areas where decisions must be made with imperfect information and the costs of type I and II errors are unequal (when the costs of both error type are equal, there are simpler tools, such as the Neyman–Pearson lemma (Neyman and Pearson, 1933), for assessing the optimal strategy). Immunology is one such area: excessive type I errors by macrophages result in infection of the host, while excessive type II errors result in autoimmune disorders (Morris, 1987).

ACKNOWLEDGMENTS

REFERENCES

Adler, J., and Wung-Wai, T. (1974). Decision-making in bacteria: chemotactic response of *Escherichia coli* to conflicting stimuli. *Science* 184, 12921294.

Apfelbach, R., Blanchard, C., Blanchard, R., Hayes, R., and McGregor, I. (2005). The effects of predator odors in mammalian prey species: a review of field and laboratory studies. *Neurosci. Biobehav. Rev.* 29, 1123–1144.

Benton, T., and Evans, M. (1998). Measuring mate choice using correlation: the effect of sampling behaviour. *Behav. Ecol. Sociobiol.* 44, 91–98.

Bialek, W., and Setayeshgar, S. (2005). Physical limits to biochemical signaling. *Proc. Natl. Acad. Sci. U.S.A.* 102, 10040–10045.

Blumstein, D. (2003). Flight-initiation distance in birds is dependent on intruder starting distance. *J. Wildl. Manage.* 64, 852–857.

Boas, M. (2006). *Mathematical Methods in the Physical Sciences*. Danvers, MA: Wiley, 864.

Broom, M., and Ruxton, G. (2005). You can run – or you can hide: optimal strategies for cryptic prey against pursuit predators. *Behav. Ecol.* 16, 534–540.

Cooper, W. (2006). Dynamic risk assessment: prey rapidly adjust flight initiation distance to changes in predator approach speed. *Ethology* 112, 858–864.

Cooper, W., and Frederick, W. (2007). Optimal flight initiation distance. *J. Theor. Biol.* 244, 59–67.

Cooper, W., and Frederick, W. (2010). Predator lethality, optimal escape behavior, and autonomy. *Behav. Ecol.* 21, 91–96.

Cooper, W., Hawlena, D., and Perez-Mellado, V. (2009). Interactive effect of starting distance and approach speed on escape behavior challenges theory. *Behav. Ecol.* 20, 542–546.

Creswell, W. (2008). Non-lethal effects of predation in birds. *Ibis* 150, 3–17.

Creswell, W., Butler, S., Whittingham, M., and Quinn, J. (2009). Very short delays prior to escape from potential predators may function efficiently as adaptive risk-assessment periods. *Behaviour* 146, 795–813.

Floreano, D., and Keller, L. (2010). Evolution of adaptive behaviour in robots by means of Darwinian selection. *PLoS Biol.* 8, e1000292. doi: 10.1371/journal.pbio.1000292

Hairston, N. G., Tinkle, D. W., and Wilbur, H. M. (1970). Natural selection and the parameters of population growth. *J. Wildl. Manage.* 34, 681–689.

Hemmi, J., and Merkle, T. (2009). High stimulus specificity characterizes antipredator habituation under natural conditions. *Proc. R. Soc. Lond. B Biol. Sci.* 276, 4381–4388.

Hemmi, J., and Pfeil, A. (2010). A multi-stage anti-predator response increases information on predation risk. *J. Exp. Biol.* 213, 1484–1489.

Janetos, A., and Cole, B. (1981). Imperfectly optimal animals. *Behav. Ecol. Sociobiol.* 9, 203–209.

Korn, H., and Faber, D. (2005). The Mauthner cell half a century later: a neurobiological model for decision-making? *Neuron* 47, 13–28.

Luttbeg, B., and Warner, R. (1999). Reproductive decision-making by female peacock wrasses: flexible versus fixed behavioral rules in variable environments. *Behav. Ecol.* 10, 666.

Marshall, J., Dornhaus, A., Franks, N., and Kovacs, T. (2006). Noise, cost and speed-accuracy trade-offs: decision-making in a decentralized system. *J. R. Soc. Interface* 3, 243.

Morris, J. (1987). Autoimmunity: a decision theory model. *J. Clin. Pathol.* 40, 210–215.

Nelson, E., Matthews, C., and Rosenheim, J. (2004). Predators reduce prey population growth by inducing changes in prey behavior. *Ecology* 85, 1853–1858.

Neyman, J., and Pearson, E. (1933). On the problem of the most efficient tests of statistical hypotheses. *Philos. Trans. R. Soc. Lond. A* 231, 289–337.

Parker, G., and Smith, J. (1990). Optimality theory in evolutionary biology. *Nature* 348, 27–33.

Perkins, T., and Swain, P. (2009). Strategies for cellular decision-making. *Mol. Syst. Biol.* 5, 326.

Rabin, M. O. (1963). Probabilistic automata. *Inform. Control* 6, 230–245.

Rao, R. (2010). Decision making under uncertainty: a neural model based on partially observable Markov decision processes. *Front. Comput. Neurosci.* 4:146. doi: 10.3389/fncom.2010.00146

Roberts, B. (1992). Neural mechanisms underlying escape behavior in fishes. *Rev. Fish Biol. Fish.* 2, 243–266.

Rock, M., Hackett, J., and Brown, D. (1981). Does the Mauthner cell conform to the criteria of the command neuron concept? *Brain Res.* 204, 21–27.

Schaik, C., van Noorwijk, M., Warsono, B., and Sutriono, E. (1983). Party size and early detection of predators in Sumatran forest primates. *Primates* 24, 211–221.

Sih, A. (1992). Prey uncertainty and the balancing of antipredator and feeing needs. *Am. Nat.* 139, 1052–1069.

Ydenberg, R., and Dill, L. (1986). The economics of fleeing from predators. *Adv. Study Behav.* 16, 229–249.

Ylonen, H., Kortet, R., Myntti, J., and Vainikka, A. (2007). Predator odor recognition and antipredatory response in fish: does the prey know the predator diel rhythm? *Acta Oecol.* 31, 1–7.

Zottoli, S., and Faber, D. (2000). The Mauthner cell: what has it taught us? *Neuroscientist* 6, 26–38.

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.