Modelling Animal Behaviour in Contests: Conventions for Resource Allocation

Matt Quinn[1] and Jason Noble[2]

[1]Centre for Computational Neuroscience and Robotics
University of Sussex, Brighton BN1 9QG, UK
[2]Informatics Research Institute, School of Computing
University of Leeds, Leeds LS2 9JT, UK
matthewq@cogs.susx.ac.uk, jasonn@comp.leeds.ac.uk

Abstract. The selective pressures affecting animal contest behaviour are investigated with an evolutionary simulation model. Two agents of differing fighting ability compete for control of an indivisible resource. Results indicate the evolution of coordinated behaviour that avoids unnecessary fighting. Detailed examination of one run shows the use of an arbitrary convention which makes ability irrelevant to contest outcome. Implications for theories of animal conflict are explored.

1 Introduction

In their quest to survive and reproduce, animals of the same species inevitably compete for resources such as food, territory, and mating opportunities. Contests may occur when two animals simultaneously attempt to gain control of a single indivisible resource. If both individuals want the resource, but only one can have it, then their genetic interests conflict, and we might expect the question of possession to be settled aggressively, through fighting. However, one of the curious facts about animal contests is how often all-out violence is avoided. From spiders [1] to elephants [2], we find that most of the time contests stop short of serious injury and are settled by what appear to be threats, signals of strength or determination, or even arbitrary conventions.

How and why do animals manage to settle contests peacefully so much of the time? Ethologists used to argue that animal contests were often non-violent because too much fighting would be bad for the species. For instance, Huxley [3] believed that animals produced and attended to ritualized threat displays in order to "reduce intra-specific damage." Unfortunately this idea fails to recognize that animals are not interested in the survival of their species, but rather in the propagation of their own genes. The difficulty is to explain why selfish individuals should be expected to show restraint in a contest. Applying game theory to evolutionary biology, Maynard Smith and Price [4] used the Hawk-Dove game to show that a tendency to stop short of all-out fighting could be explained in terms of evolved self-interest. Given the reasonable assumption that the contested resource is valuable but not worth risking serious injury for, the

Hawk-Dove game demonstrates that aggressive and peaceful strategies will both do well when they are rare in a population, which means that in the long run we should expect a stable state in which both are present.

Communication is one way of minimizing violence in contests: in theory, if both competitors exchange information about their fighting ability, and then the weaker defers to the stronger, most fights are avoided. (Equivalent mechanisms might be based on other relevant variables, such as the subjective value of the resource for each animal.) Game-theoretically inclined biologists have debated at length whether such a signalling system would be evolutionarily stable. There is widespread agreement that if information about the opponent's fighting ability cannot be faked, as when one animal is simply much bigger than the other, the information will be attended to. Much more problematic is the case where a signal indicating strength could in theory be given by a weak animal, i.e., it is possible for weak competitors to bluff or cheat. One school of thought [5] says that there would be nothing to prevent the evolution of bluffing in such cases, and that the signalling system would therefore collapse—once bluffing becomes prevalent, there is no longer any point in attending to your opponent's signal. An opposing view [6] suggests that under certain circumstances it is not worthwhile for a weak animal to signal dishonestly, as bluffs might be challenged by stronger competitors to disastrous effect. All parties agree that honest signalling, if it is to occur, needs to be stabilized by costs of some sort—it must not be worthwhile to bluff. A number of cost regimes that could conceivably underlie a signalling system advertising fighting ability have been identified [7], e.g., a vulnerability handicap, whereby stronger animals place themselves in a position where they are vulnerable to attack, a risk they could not afford to take if they were weak. However, due to the difficulties involved in measuring the costs, risks, and benefits of contests empirically, it is still an open question as to whether animals really use signalling systems to minimize violent conflicts.

We wanted to find out more about the selective pressures impinging upon animal contest behaviour, and have therefore constructed an artificial life model of contests between two agents over an indivisible resource. In our model the agents are controlled by recurrent neural networks which are shaped over generational time by a genetic algorithm. The world is two-dimensional and complies with basic physics such as inertia and friction. Contests take place in a circular arena; the agents can fight by ramming each other, and must physically occupy the centrally located resource in order to win possession of it. Agents differ in their fighting ability, and although they can sense their own ability, they cannot directly detect the ability of their opponent. By allowing behavioural strategies to evolve over time in this model, we hoped to answer various questions about contest behaviour. For example, should agents behave differently depending on their fighting ability? If the opponent makes an aggressive approach, should an agent take the threat seriously? When a contest is settled peacefully, how do the two agents reach a decision as to who gets the resource and who backs down? Will it be evolutionarily stable for the agents to exchange truthful signals as to their fighting ability?

One of our aims in this paper is to interpret the evolved strategies in the light of current biological theory. If a particular theory is useful in understanding the evolved behaviour of an agent controlled by a recurrent neural net, that should increase our confidence that it will be useful in understanding the behaviour of real animals. At the same time, we believe that our semi-realistic simulation may suggest novel ways of thinking about contest behaviour. Most of the biological ideas discussed so far are based on insights from game-theoretic models, and, for reasons of mathematical tractability, such models reflect simple situations with a very small range of strategies, and, at best, a minimal treatment of space and time. For instance, in a model suggesting honest signalling of strength by weaker competitors [6], all individuals are either weak or strong, and they make only two decisions: which of two arbitrary signals to make, and then whether to flee, attack, or attack after a pause. In our model, contests are played out over space and time, introducing possibilities such as chases and stand-offs. The need to ram the other agent in order to cause damage means that aggressive intentions can be perceived; simulating momentum automatically implements a notion of commitment to an attack, as an agent may be going so fast that it could not possibly slow down before impact. Factors like these may well lead to the evolution of strategies that cannot be expressed in a simple game-theoretic model, and the same factors may be critical to an understanding of behaviour in real animal contests. One of the authors [8] has previously looked at the evolution of contest behaviour in a richer context than that afforded by paper-and-pencil game theory, and found that honest signals of fighting ability were unlikely to evolve. However, the simulation environment in this earlier work was only of intermediate complexity.

Another goal for the work is to illustrate the potential for artificial life models to examine the evolution of communication from genuinely non-communicative origins. Movement is the only way that one agent can impinge upon the sensory world of the other, and if the agents come to exchange signals of fighting ability, these signals will have to develop from movements that do not originally have signalling as their function. We want to move away from the provision of arbitrary signalling channels, as found in most artificial life models of the evolution of communication. If you start with agents that already have the ability to send and receive signals on a dedicated channel, then you cannot be investigating the *origin* of communication, only the parameters under which it might be evolutionarily stable once established (see Quinn, this volume, for further discussion of this point).

2 The simulation

Contests commence with two agents placed at random positions[1] and orientations in a simulated arena 80cm across (Fig. 1). The circular body of each agent is 10cm in diameter. The contested resource is also 10cm across and is located

[1] Agents must start at least 5cm away from the arena edge, 5cm away from the resource, and 10cm away from each other.

in the centre of the arena. Agents have a fighting ability drawn from the range {0.1, 0.3, 0.5, 0.7, 0.9}, and an energy level that starts at one and is reduced through damage; an energy level of zero indicates serious injury. Agents can accelerate to a maximum speed of approximately 15cm/s. Time is simulated in slices of 0.1s duration. Inertia, friction and elastic collisions are implemented in the simulation.

- 7 distance sensors (d0-d6) for detecting the opponent.
- 2 sensors (Rl, Rr) for detecting the resource.
- An omnidirectional distance sensor (Rd) and a contact sensor (Ro) for detecting the resource.
- A bumper sensor (bp) for detecting collisions.
- Sensors for fighting ability (A) and energy level (E).

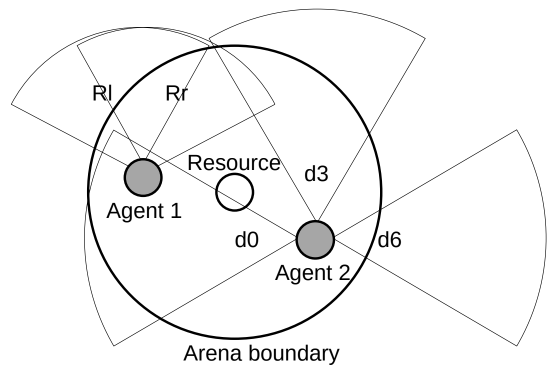

Fig. 1. A list of the agents' sensory inputs, plus a diagram showing the arena, the resource and the competing agents, to scale. Both agents are facing north. Agent 1 illustrates the arrangement of the two resource sensors, Rl and Rr. Note that agent 1 cannot see the resource from the position shown. Agent 2 illustrates 3 of the 7 proximity sensors for detecting the opponent. The sensors are equally spaced around the front half of the agent's body; in the interests of clarity, sensors d1, d2, d4 and d5 are not shown. Note that agent 2 can detect agent 1 with d0 (and also with d1, not shown).

Agents can inflict damage by ramming each other. The agents' bodies are equally vulnerable at all points, but in a collision A will only damage B if A is moving towards B at the time. Thus, in a head-on collision, damage is mutual, but if A rams B while B is moving away or stationary, only B will be damaged. Damage is subtracted from the target's energy level, and is equal to $K \times F_{\text{attacker}} \times (1 - F_{\text{target}}) \times \frac{\text{speed}}{\text{max. speed}}$; K is a constant set equal to 4, and F is fighting ability. The essence of this is that fighting ability is valuable for both attacking and defending, and high speed means more damage. Fighting is a risky proposition: for example, if two agents of only average ability (0.5) ram each other at full speed, they will inflict 1 unit of damage, meaning serious injuries to both. Once an agent has suffered a serious injury, it is out of action for that contest and disappears immediately from the arena.

Agents can run away instead of fighting. The boundary of the arena is not a wall, and if an agent crosses the boundary it is assumed to have fled the contest, and will disappear from its opponent's sensors. (Note that this does not mean that the agent remaining in the arena automatically gains the resource.)

A contest will normally end when one contestant gains possession of the resource, which an agent achieves by keeping the centre of its body continuously within the resource zone for 10s. Contests can also end if both agents flee the arena, or if a 30s time limit expires. (If an agent is within the resource zone at the end of 30s, the contest is extended for up to 10s to give the agent a chance to gain possession.)

Agents have 14 sensory inputs (detailed in Fig. 1) allowing them to perceive each other, the resource, and their internal state. Agents are controlled by arbitrarily recurrent artificial neural networks. The thresholds, weights and decay parameters, and the size and connectivity of the network are all genetically determined. Agents are equipped with left and right motors, much like a Khepera robot, and four binary outputs from the neural network control forward and backward drives for each motor (if the forward and backward drives for a motor are both activated, output is zero).

Each generation, an agent participates in 25 contests, one for each of the possible combinations of their own and their opponent's fighting ability. Agents also participate in a further five trials in which there is no opponent present in the arena; these trials were included in order to ensure selection pressure for the ability to take an uncontested resource.

In terms of fitness, gaining control of the resource is worth 100 units. Injury is a fitness cost that scales with the damage an agent has suffered, to a maximum of 200 units for a serious injury. There is also a modest benefit for fleeing a contest promptly, scaled according to the time remaining when the agent flees, up to a maximum of 10 units for instantaneous flight. These costs and benefits have been chosen such that the resource is worth having, but not worth risking serious injury for. Previous models [5, 6, 8] have shown that this balance between resource benefit and injury cost poses the most interesting strategic problem for the contestants. (If the resource is made too valuable, competitors will fight to the death. If the resource is not valuable enough, competitors will not be tempted to fight at all.) The inclusion of a benefit for leaving the contest early captures the idea that participating in a contest always entails an opportunity cost. The costs of movement, including attacks, are assumed to be negligible: this reflects empirical work on spider contests [1] showing that the long-term fitness costs of serious injury—and, of course, death—are orders of magnitude greater than other costs such as energetic expenditure associated with threat displays.

The evolutionary engine for the simulation was a genetic algorithm operating on a population of 200 agents. Selection was fitness-proportional, based on mean fitness per contest, with the elite individual retained unchanged in the next generation. Crossover was used 75% of the time. The agents' neural networks were represented in the genotype using a topological encoding scheme, which necessitated several different types of mutation. The rate of non-structural mutations (e.g., weight changes) was 1.0 per genotype, the probabilities for adding or deleting a neuron were 0.0025 and 0.02 respectively, and the probabilities for adding, deleting and reconnecting a synapse were 0.01, 0.05, and 0.03.

3 Results and interpretation

Ten runs of at least 20000 generations were conducted, each with a different random seed value. Initial controllers consisted of three neurons with randomly determined connections and parameters. We will first summarize general findings across the ten runs, and then examine a single run (run zero) in greater detail.

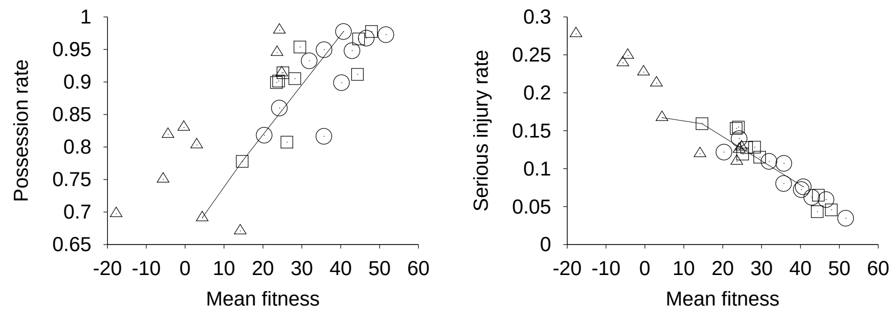

Fig. 2. Left: proportion of contests in which an agent takes possession of the resource, plotted against mean fitness at three points in time across ten runs. **Right:** rate of serious injury (per agent per contest) by mean fitness score, at three points in time across ten runs. **Both:** in these and subsequent figures, the mean state of a population between generations 200–400, 5000–5200, and 18000–18200 is represented with a triangle, a square and a circle respectively. The relevant points for run zero are joined with a line to show progress over time.

One question we can ask is whether or not the agents managed to settle contests with a minimum of violence, as animals do. Figure 2 (left) shows the proportion of contests in which one agent or the other took possession of the resource, plotted against mean fitness scores. Focusing on the situation after 18000 generations (points marked with circles) we can see that the possession rate is at least 80% in all ten runs. This indicates that the agents are at least managing to allocate the resource most of the time, rather than, for example, both fleeing, or becoming locked in a stalemate. Note also that possession rates and fitness both tend to increase over generational time. Figure 2 (right) shows that the prevalence of serious injury is negatively linked to fitness, that it decreases over time, and that by 18000 generations it is as low as around 5% in some runs. If the agents' method for resolving contests was to fight until one or the other was seriously injured, the rate would be 50% (or even higher given that sometimes both would be injured simultaneously). Thus, rates of 5–10% indicate that a more peaceful system has evolved.

This last point is reinforced if we consider the simple fact that fitness scores after 5000 generations are all positive. In a perfect world, agents might toss a coin and allocate the resource accordingly, without a fight. Each agent could therefore expect the best possible mean fitness[2] of $100/2 = 50$. If fighting is occurring, injuries will mean that fitness scores are lower—for example, if agents were to inflict an average of 50% damage on each other, then mean fitness would fall to -50. So the observation that fitness scores are greater than zero suggests that fighting has been minimized.

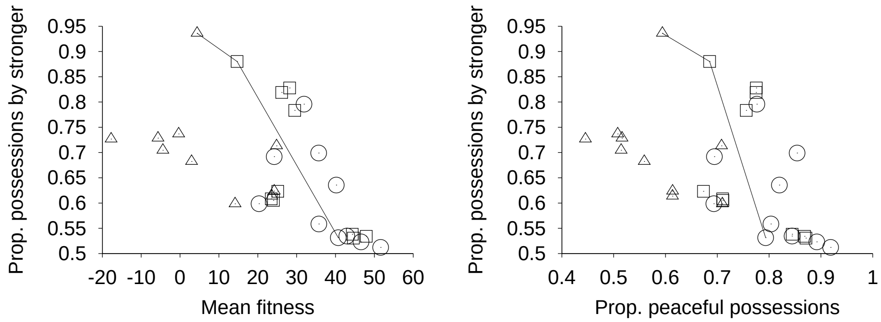

Fig. 3. Left: proportion of possession events in which the stronger agent takes the resource, plotted against mean fitness at three points in time across ten runs. **Right:** proportion of possession events in which the stronger agent takes the resource, plotted against the proportion of peaceful possessions (i.e., one agent takes the resource without either agent being damaged), at three points in time across ten runs. See Figure 2 for an explanation of symbols used.

Could the agents be signalling their fighting ability to each other, with the weaker deferring to the stronger? Figure 3 suggests that this is not the case for most runs. If such a signalling system had evolved, the stronger of the two agents should usually take the resource without a fight. Figure 3 (left) shows that, early on in a run (points marked with triangles) the stronger agent wins about 70% of the time. But as evolution progresses, this figure falls, and in about half of the runs it falls close to the chance level of 50%. Furthermore, Figure 3 (right) shows that high rates of possession by the stronger agent are associated with fewer peaceful possessions, i.e., that when victory is to the strong it is not without a fight. The general profile that emerges from late in the runs (points marked with circles) is one of peaceful allocation of the resource, high mean fitness scores, and fighting ability not being an important predictor of success. The latter point does not match what we would expect if a signalling system was in place.

If the agents have evolved a way of allocating the resource without violence, and if fighting ability has become unimportant, then what might be happening? We extended run zero to 68000 generations and looked at the results in greater

[2] In fact the true theoretical optimum in our simulation is 62.5 due to the bonus for early fleeing and the 5 trials out of 30 conducted with no opponent.

detail in order to find out—Figures 2 and 3 show that, at 18000 generations, run zero is reasonably representative of the ten runs.

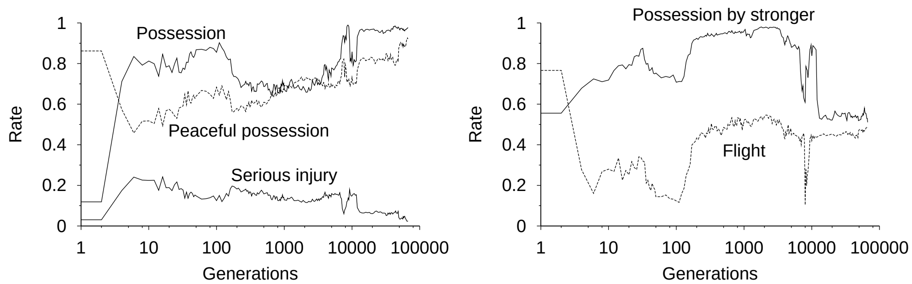

Fig. 4. Left: proportion of contests ending in possession of the resource, proportion of peaceful possessions, and rate of serious injury, plotted over generations for run zero. **Right:** proportion of possessions by stronger agent, and rate at which agents fled the arena, plotted over generations for run zero. Note that the time axes are log-scaled.

Figure 4 shows various statistics for run zero. Unsurprisingly, given their randomly connected networks, agents in the initial generation were not likely to gain possession of the resource; they tended to move in a random direction and thus flee the arena without fighting. However, we see that the ability to move inwards and gain possession of the resource develops very quickly, and by about generation 8 we see the highest levels of serious injury and the lowest levels of peaceful possession as both agents try to move onto the resource. Between generation 10 and generation 75, contests become more peaceful and the stronger agent becomes more likely to win. This is because the agents have started to use their fighting ability sensor, and will move in more confidently if they are strong, and may flee immediately when very weak. By generation 100 we see a slow, clockwise spiralling motion in towards the resource; the bias towards right turns is maintained throughout the run. As the generations pass the behaviour of the agents becomes more and more strongly modulated by their fighting ability. Around generation 8000 there is a spike in the possession rate, and a drop in the rate at which the stronger agent wins. This occurs because the agents adopt a slow, waggling motion, and become very reluctant to collide with each other. Around generation 9000 fast, smooth motion is recovered. Typically, both agents orbit the resource zone, and each avoids the other, but the stronger is prepared to turn in more tightly, and thus gains the resource. Then, around generation 12200, fighting ability rapidly becomes less useful as a predictor of victory. The agents circle each other, and one seems to get a positional advantage on the other, irrespective of being weak or strong. The first agent chases the second out of the arena, and then returns for the resource. Given this behaviour, the rate of serious injury drops, and the rate of peaceful possessions jumps to 80% or more. The pattern is maintained for the rest of the run.

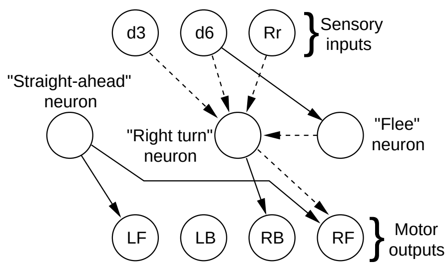

Fig. 5. Schematic diagram of an evolved neural network from generation 17,900 in run zero. Excitatory and inhibitory connections are indicated by solid and broken lines respectively. Neurons and connections that are non-functional or relatively unimportant are not shown.

This behaviour appears to be a convention for settling most contests without violence. We analyzed one of the evolved neural nets in order to see how the agents managed to do this (Fig. 5). The evolved controller represents a balance between two tendencies: a spiralling taxis towards the resource, and movement relative to the opponent depending on orientation. The "straight-ahead" neuron is alway firing, due to a negative threshold and connections to various inputs; it tries to drive both forward motors. However, in the absence of significant sensory input, the "right turn" neuron dominates, switching on the right-backwards motor and switching off the right-forwards motor, resulting in rotation to the right. When the resource is detected with the Rr sensor, the right turn neuron is inhibited, and the agent moves forward. When the resource is no longer visible, the agent turns right again—this is enough to implement a spiralling movement in towards the resource.

Detecting the opponent with either the d3 (straight ahead) or the d6 (90º right) sensor inhibits right turns and causes forward movement. This will have two quite different results: if the opponent is detected dead-ahead, the agent will advance aggressively. If the opponent is detected to the right side, the agent will advance but this will ultimately lead to its running away, chased by the opponent. This happens because the d6 sensor also activates the "flee" neuron, which has a very slow decay rate, and functions as a switch. Once the flee neuron is firing, it permanently inhibits the right turn neuron, and so the agent keeps moving in a straight line, which will always take it out of the arena. The opponent, having a very similar neural net, will advance aggressively, "chasing" the fleeing agent from the arena, and then return to the resource when its right-turning taxis regains control.

Interestingly, the evolved networks use the random placement of the agents in the arena as an arbitrary way of settling the contest without a fight. Both agents will turn right from their initial positions, the lucky one will catch sight of the other dead ahead, and the unlucky one will first notice the opponent to its right. The unlucky agent will be chased away; the lucky one will take the resource uncontested. The convention is not quite arbitrary, however: it appears to have evolved because having another agent approach you from the side does indeed put you at a disadvantage, given that you would need to turn and build up speed again before you could attack that agent.

4 Conclusion

Our simulation has shown that a population of selfish agents can evolve relatively peaceful strategies for resource allocation. The agents exploit a convention that is, on the one hand, rooted in facts about their sensory and motor capacities, and, on the other, rendered arbitrary because of random starting positions. We had not anticipated the evolution of a convention for avoiding conflict: the evolved agents found a third way between the rival game-theoretic predictions of honest signalling [6] and poker-faced inscrutability [5]. Arbitrary conventions have been discussed in the animal-contest literature before, but were generally not thought to be relevant to situations where the competitors differed in fighting ability. Our model suggests that it may be worth looking for such conventions in a wider range of animal contest situations. More generally, we believe that our findings illustrate the value of an open-ended approach to modelling animal communication and interaction.

Acknowledgements

We acknowledge computing support from the Sussex High Performance Computing Initiative.

References

1. Riechert, S.E.: Spider interaction strategies: Communication vs. coercion. In Witt, P.N., Rovner, J., eds.: Spider Communication: Mechanisms and Ecological Significance. Princeton University Press (1982) 281–315
2. Poole, J.H.: Announcing intent: The aggressive state of musth in African elephants. Animal Behaviour **37** (1989) 140–152
3. Huxley, J.S.: Ritualization of behaviour in animals and men. Philosophical Transactions of the Royal Society of London: Biological Sciences **251** (1966) 249–271
4. Maynard Smith, J., Price, G.R.: The logic of animal conflict. Nature **246** (1973) 15–18
5. Caryl, P.G.: Communication by agonistic displays: What can games theory contribute to ethology? Behaviour **67** (1979) 136–169
6. Hurd, P.L.: Is signalling of fighting ability costlier for weaker individuals? Journal of Theoretical Biology **184** (1997) 83–88
7. Vehrencamp, S.L.: Handicap, index, and conventional signal elements of bird song. In Espmark, Y., et al., eds.: Animal Signals: Signalling and Signal Design in Animal Communication. Tapir Academic Press (2000) 277–300
8. Noble, J.: Talk is cheap: Evolved strategies for communication and action in asymmetrical animal contests. In Meyer, J.A., et al., eds.: Proc. Sixth International Conference on Simulation of Adaptive Behavior, MIT Press (2000) 481–490