

An Agent-based Model for Pain Expression

COMPLEX MRES SUMMER PROJECT

Elizabeth Gallagher

Supervisors: Dr Peter Bentley and Dr Amanda Williams

August 22, 2013

Abstract

Pain motivates escape, healing and recuperation. If pain can be effectively communicated in the presence of caregivers then these objectives can be achieved sooner. These aims can further be aided if pain behaviour can be suppressed in the presence of antagonists, where vulnerability can be taken advantage of. Hence, in certain social situations the amount of tissue damage can be uncorrelated to the degree of pain behaviour.

By considering the interactions of heterogeneous individuals in a range of different environments the conditions for which pain expression is favourable can be established. This project aims to achieve such a goal by developing an agent-based model where agents can either express or suppress pain if randomly injured, and aid or ignore the pain expression of others. Interactions between agents will incur costs and benefits, these will affect which agents are chosen for interactions, which agents will “die” and which agents will be selected for parenthood. Hence, the proportions of strategies in an agent population found using different parameter values will reflect the conditions needed for pain expression to prevail.

The results found are intuitive and successfully simulate the co-occurrence of pain expression with its ability to be detected and helped by others. Furthermore, results predict that social support motivates expression and an increase in social threat provokes less expression; results which have been observed experimentally.

Contents

1	Introduction	1
2	Background	3
2.1	Pain behaviour	3
2.1.1	Pain expression and its interpretation: costs and benefits	4
2.1.2	Social influence on pain	5
2.1.3	An evolutionary perspective of social pain behaviour	7
2.2	Agent-based models	8
3	An agent-based model for pain expression	10
3.1	Agents	10
3.2	Pain and interactions	14
3.3	Replacement	15
3.4	Parameters	15
4	Experiments and results	21
4.1	Preliminary experiments	21
4.1.1	The initial agent population	21
4.1.2	Simulation variability	22
4.1.3	The number of iterations	24
4.1.4	The number of repeated simulations	24
4.2	Cost and benefit effects	25
4.3	Danger	28
4.4	Sociability	29
4.5	Antagonists	30
4.6	Inheriting connectedness	30
5	Discussion	33
5.1	General observations	33
5.2	Cost and benefit parameters	34

5.3	Danger	36
5.4	Sociability	36
5.5	Antagonists	37
5.6	Inheriting connectedness	37
6	Conclusions and extensions	39
	Appendix	41
	References	69

1 Introduction

Pain experience and pain behaviour are distinct entities with their own functions and influences (Sullivan *et al*, 2006).

Pain signals to an organism that a situation is potentially harmful, the organism's behaviour in response to this aims to escape or reduce pain and also to maximise the chance of healing if there is an injury. The adaptive advantage of feeling pain is of an extent that humans who have never been able to experience pain have a significantly shorter lifespan than those who have (Damasio, 1999).

The functions of pain behaviour can be categorised as protective (e.g. guarding, rubbing) or communicative (e.g. facial grimacing, vocalisations) (Williams & Craig, 2006). Although both of these behaviours can be observed, and therefore both can result in empathy and aid from caregivers, the former also allows an individual to avoid further damage and recuperate, and the latter can be manipulated (Sullivan *et al*, 2006). Observed pain behaviour does not always correlate to the amount of pain an animal is in, and often the social situation will have an effect on the extent of the behaviour (Williams, 2002).

Pain behaviour comes at a physical energy cost, and also a cost due to it drawing attention to vulnerability (Williams, 2002). The benefits of communicative pain behaviour are dependent on there being others around which can aid the expressing individual. However, it is not always the case that others will decide to help, as in some circumstances the cost of helping overrides its benefit. The benefits of helping an expressing individual can occur when there is a chance of reciprocation (Trivers, 1971), due to kin selection or due to group selection (Nowak, 2006). Hence, the influences towards whether expression is beneficial or not require careful consideration of many factors and outcomes of certain interactions.

Despite the fact that these cost and benefits indicate that expressive behaviour is likely to be strategic, an evolutionary viewpoint of the study of pain behaviour has generally been overlooked (Williams, 2002). However, an evolutionary approach could give rise to a better understanding of the processes and social conditions leading to pain behaviour. This project aims to undertake such an approach.

An agent-based approach to modelling allows a model to be made where the net effect of the interactions of heterogeneous individuals can be seen, sometimes revealing and replicating complex phenomena. Hence, in this project an agent-based model will be developed in order to explore the circumstances for which pain expression can persist or decline in different social environments. The agents of this model will have either an expression or suppression strategy when injured, and will either be altruistic or selfish when interacting with an injured agent. Interactions will incur costs and benefits to both the agent in pain and the interacting agent based on their given strategies. Varying these cost and benefits, and other parameters, we can simulate different environments and find how these effect the occurrence of pain expression.

A short review of the pain behaviour literature, specifically focusing on the effect of social influences, and a general agent-based modelling introduction will be given in Section 2. Section 3 will describe the agent-based model developed in this project. Preliminary experiments and also the influence of various parameter changes on the degree of expression will be considered in Section 4 and the results from these will be discussed in Section 5. These parameter changes will include; the cost and benefit parameters, how dangerous the environment is, how sociable the individuals are, and what can be inherited from one generation to the next. The effect of an antagonist influence will also be considered. Finally, Section 6 will conclude findings and also discuss extensions to the model.

2 Background

2.1 Pain behaviour

Pain behaviour has been studied in many animals, including mice (Langford *et al*, 2010), cows (Barrier *et al*, 2012), fish (Black, 2009) and humans (Ranger, Johnston & Anand, 2007; Peeters & Vlaeyen, 2011).

Wall (1979) suggested that pain is more important for healing than injury avoidance as often pain intensity is weakly related to the amount of tissue damage. He describes three stages of behaviour after injury, first is the immediate stage where escaping, fighting and obtaining aid takes priority. Next is the transitional phase where the organism feels pain, and is agitated and aggressive. Finally, a recuperation phase begins, which is characterised by more pain, sleep and inactivity. Alternatively Bolles and Fanselow (1980) have suggested that after injury animals first perceive their environment. This teaches the animal to fear the source and the context of the injury. Hence, they suggest that defensive behaviour after injury (such as analgesia, freezing and fighting) is a response to a memory of fear, rather than the pain itself.

It is important for the benefit of animal welfare research and for clinicians using analgesics to be able to interpret when an individual is in pain. Visual scales of pain expression have been made to aid this interpretation in nonverbal animals (see Figure 1), otherwise pain can be communicated

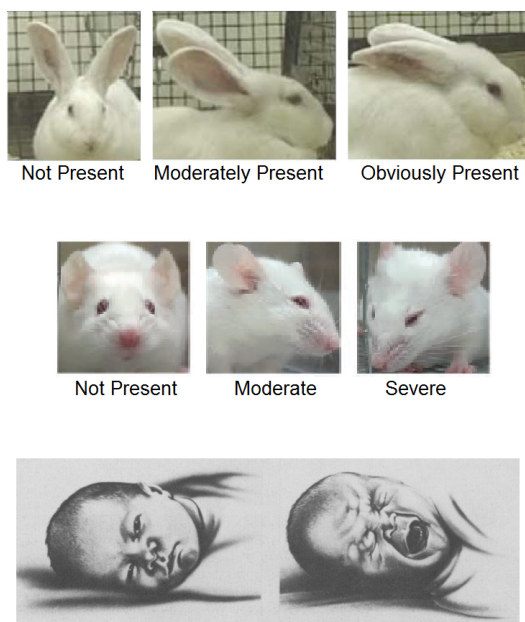


Figure 1: Examples of judging the extent of pain from nonverbal expression. Reproduced from Langford *et al* (2010) (above), Keating *et al* (2012) (center) and Grunau & Craig (1987) (below).

verbally. However, in some cases behavioural cues can also reflect states of stress. This is the case in infants, where crying could indicate agitation rather than pain (Ranger, Johnston & Anand, 2007).

Fordyce (1976) argued that pain expression developed and is maintained by reinforcement, where the amount of pain expression is dependent on its reward value. Hence, if a particular behaviour elicited positive responses from others (e.g. by attention or pain relief) then there would be a higher probability of these behaviours being repeated in the future. This operant theory supports an idea that aid can be taken advantage of because of social reinforcement. Hence, an amplification of pain behaviour does not necessarily mean an increase in pain.

2.1.1 Pain expression and its interpretation: costs and benefits

Attention and aid can be given when pain behaviour is observed by another individual. However, in situations of social threat pain behaviour can show vulnerability and lead to weakness being taken advantage of. This could come in the form of a predator, a conspecific taking social advantage or an individual stealing resources, such as food.

By not expressing pain an individual can maintain their resources. However, in social groups being dishonest about capabilities could lead to risks, for example being left behind in a group hunting situation because an injured individual was assumed to be healthy (Williams, A. C. D. C., personal communication by email, August 14, 2013).

Communicative pain expression can only be beneficial in an environment where others are able to recognise it and help, and this comes with its own costs and benefits. By being able to interpret pain behaviour in other individuals there can be a reciprocity benefit if help is given (Trivers, 1971), a benefit if a relation is helped (Nowak, 2006) and also the possibility of avoiding danger yourself. However, there is an energy cost to helping an individual in pain and it could also put the helping individual into a dangerous situation.

Thus, pain expression can be seen to vary with social context. In the following section we will describe some of the studies into these social situations.

2.1.2 Social influence on pain

Research into the impact of social factors on pain behaviour and the verbal report of pain has mostly looked at the effects of social support and social threat (Wiech & Tracey, 2013). These can illustrate some of the benefits and costs of expressing pain described above.

Social support

Pain behaviour has been shown to vary in situations where there is social support. In a study by Brown *et al* (2003) it was found that individuals reported less pain if they were with another person than if they were alone. Additionally, in a more recent study (Coan, Schaefer & Davidson, 2006) it was found that individuals awaiting a pain stimulation who were holding the hand of their spouse showed less threat-related activation in certain neural regions than those who were holding the hand of a stranger or not holding a hand at all. Furthermore, this effect was stronger in the couples which rated their marriage quality highly.

In a study by Block, Kremer & Gaylor (1980) chronic pain patients were asked to report their pain level whilst being observed by either their spouse or by a neutral observer. Patients who reported that their spouse was particularly attentive in responding to their pain behaviour reported higher levels of pain when being observed by their spouse than by the neutral observer. Conversely, those patients which reported that their spouse was unattentive, reported lower levels of pain than when being observed by the neutral observer.

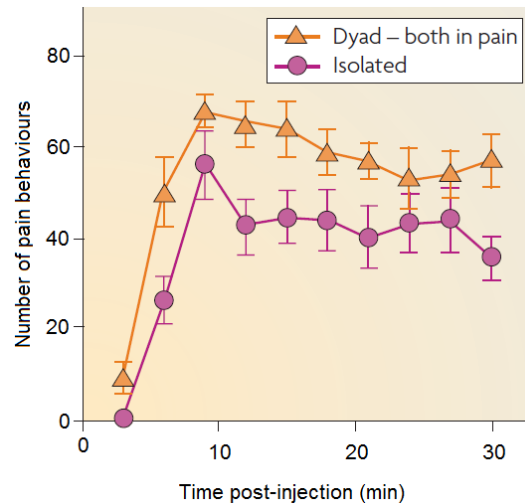


Figure 2: Pain behaviours after a noxious injection of mice in pairs (dyads) or isolated. Adapted from Langford *et al* (2006).

Additionally, Langford *et al* (2006) have reported increased pain behaviours (specifically writhing behaviour) when two mice which were cage mates were both injected with a noxious stimulus and were able to observe one another, than in isolated mice. Furthermore, the writhing behaviour in these pairs co-occurred in time, at levels significantly higher than by chance, and more so between cage mates than between strangers (see Figure 2). Langford *et al* (2006) suggested that these findings implied that there was pain communication between mice, resulting in pain empathy.

Social threat

There are cases of reduced pain behaviour when an animal is faced with a situation with an increased social threat, this is generally explained as being caused by stress-induced analgesia. If further injury is threatened, pain behaviour would not be advantageous to an organism's survival, and furthermore predators have a preference for injured prey (Butler & Finn, 2009). Mogil (2009) commented that the reason it can be difficult to recognise pain in rodents is because they are prey animals, and thus predators would see an easy target if they showed pain too expressively.

This has been reported experimentally by MacIntyre *et al* (2007) where the presence of observing humans was found to stop pain-related behaviour in male mice. Moreover, Langford *et al* (2011) have found that when male mice in pain can observe a stranger male mouse, then they display less pain behaviour than when tested alone. They speculated that this was a form of stress-induced analgesia where pain behaviour is inhibited in order to aid the possibility of escape. However, they also found that when there was only partial contact between similar male dyads mice displayed more pain behaviour. They proposed that this was because pain sensitivity varies with social threat severity.

In a report by Peeters and Vlaeyen (2011) humans were led to believe they would either be receiving 5 pain units, between 5 and 20 pain units or between 1 and 5 pain units from another person (these three settings being the correlates to social threat). During the pain administration the participants actually always received 5 pain units. They were filmed and after asked how much pain they thought they had received. The results showed that an increase in the perceived social threat led to a decrease in nonverbal pain expression,

but an increase in the patients' verbal report of pain severity (see Figure 3).

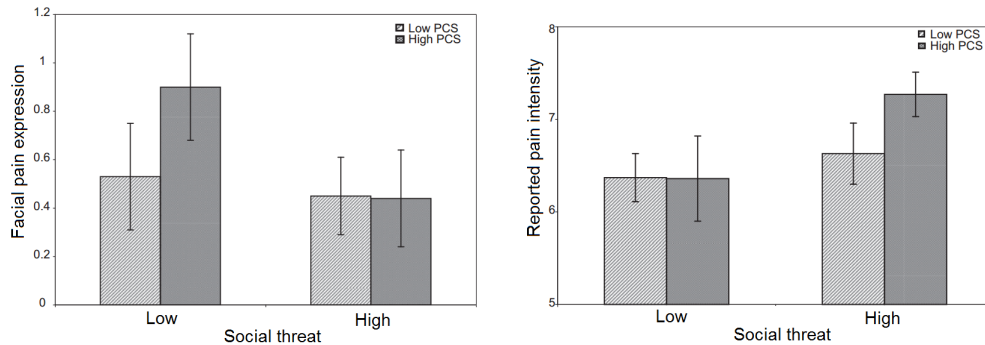


Figure 3: Facial expression (left) and reported pain intensity (right) in high and low instances of perceived social threat. Where participants either often (dark) or not that often (light) experience pain as threatening. Adapted from Peeters & Vlaeyen (2011).

The examples above show a suppression of pain behaviour when there is a social threat, a response in order to avoid further pain and to facilitate escape in a possibly dangerous situation (Langford *et al*, 2011).

Hence, pain behaviour has shown to be recognised by conspecifics and in situations of social support pain behaviour can be amplified and in situations of social threat it can be suppressed.

2.1.3 An evolutionary perspective of social pain behaviour

The idea that pain expression is maintained by reinforcement (i.e. Fordyce's operant mode (1976) discussed in Section 2.1.2) does not always hold true. Under this model the ability to suppress pain expression would be unrewarded, but in certain circumstances there may be a survival advantage (such as when faced with an antagonist) (Williams, 2002). Furthermore, the functions of pain behaviours are more complex than simply immediate avoidance and recuperation, and it is been shown that the social environment also has an impact. Hence, Williams (2002) comments that pain behaviour should be looked at from an evolutionary perspective, thus providing a broader understanding of the mechanisms

behind pain expression.

From the evolutionary perspective it is suggested that both expressing pain in the presence of caregivers and its recognition by others have co-evolved (Williams, 2002). Additionally, it assumes that pain behaviour has a communicative function which can be used to promote empathy and support from others, and in situations of social threat it can also be suppressed.

Modelling a dynamic social environment in which individuals can either express pain to others or not could lead to interesting conclusions when determining the circumstances for which pain expression exists. In such a model it would be necessary to consider a population of autonomous individuals who either expressed pain or not if injured and furthermore would help an individual in pain or not. Thus, agent-based modelling is an appropriate paradigm for such a task.

2.2 Agent-based models

Agent-based modelling is a computational method in which decision-makers (or agents) interact through predefined rules (Farmer & Foley, 2009). Unlike other approaches, agent-based simulation enables models to be built where individual heterogeneity, agents' decision rules, adaptation and learning can be represented (Gilbert, 2008). This approach can simulate the emergence of the dynamics of a whole system just from the interactions on an individual scale, and this can sometimes reveal quite complex behaviour. Furthermore, agent-based models allow *in silico* experiments to be run which in real life would be difficult or impossible to conduct. Thus, agent-based modelling has been considered as another approach to scientific research (Axelrod, 1997).

The characteristics of agents have been described by Macal & North (2006) as follows:

1. Agents are identifiable, discrete and self-contained. They have a set of characteristics and rules for behaviour.
2. Agents live in an environment in which interactions with other agents happen. Agents can respond to this environment and interact with and recognise other agents.
3. Agents can have goals to achieve.

4. Agents can function independently of their environment and other agents.
5. Agents can learn and change their behaviours based on experience (which requires memory).

Learning can be simulated in both the individual agents and in the population. It can be modelled in three ways: agents learning from their own experience; evolutionary learning, where dead agents can be replaced by better agents; and social learning, where agents imitate others. The environment an agent is in may have no effect on the agent, or may include barriers for agent movement or the influence of resource depletion or crowding.

Agent-based models have recently become a popular tool in the life sciences, ecology and social sciences (Niazi, & Hussain, 2011). In the social sciences, the earliest agent-based model was the simple dynamic model of segregation by Schelling (1971). Pennies and dimes were moved around a grid, the coins were tolerant of neighbours from a different type, but preferred neighbours of their own type. By moving “discontent” coins, whose neighbours were more than half not of their type, to another neighbourhood, segregation can be seen (Figure 4).

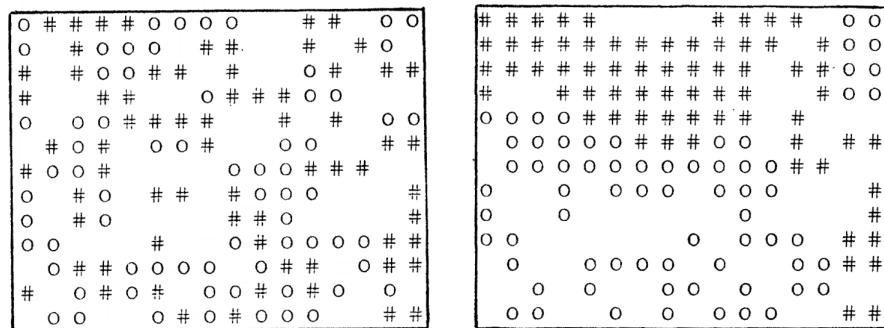


Figure 4: Initial random distribution of pennies and dimes (left) and segregation after movement (right). Adapted from Schelling (1971).

Other work has included simulating the behaviour of bumble bees on the comb (Hogeweg & Hesper, 1983), flocking behaviour (Reynolds, 1987), the evolution of cooperation (Axelrod & Dion, 1988) and whole artificial societies (Epstein & Axtell, 1996).

3 An agent-based model for pain expression

The model developed as part of this work aims to explore the circumstances for which pain expression can persist or decline in different social environments.

This model will consist of agents foraging, randomly becoming injured, occasionally interacting with one another and dying. For simplicity agents will interact in a nonspatial environment. Every iteration a number of agents will randomly become “injured”, and therefore be in “pain”. Interactions can occur between an agent in pain and any other agent. Interactions incur costs and benefits, where high costs can mean that agents playing certain strategies can die or be unlikely to be selected for parenthood. Hence, as in natural selection, agents with the most beneficial strategies for a particular environment will be the most successful and the number of those with other strategies will decline.

An outline of the model is shown in Figure 5 and agent components are summarised in Figure 6. Before being implemented, each of the model components will be described in detail. The MATLAB code used can be found in Appendix K.

3.1 Agents

The components which make up an agent are as follows: an energy level, age, a connect- edness score, an expression strategy, an altruism strategy and a time-out score. The two strategy components are unaltered in an agent’s lifetime, however the others can change. Alternatively, goal-directed agents could be used where agents could update their strate- gies in order to maximise their energy (which is equivalent to fitness as this determines the likelihood of parenthood).

Each of the components of an agent will be described below.

Energy and age levels

Agents live for a number of iterations or until their energy reaches zero, at this point the agent will “die” and be replaced by another. The initial ages of agents are taken from a random uniform distribution between one and a maximum age. Agents are all given the same energy level to begin with and at each iteration uninjured agents “forage”

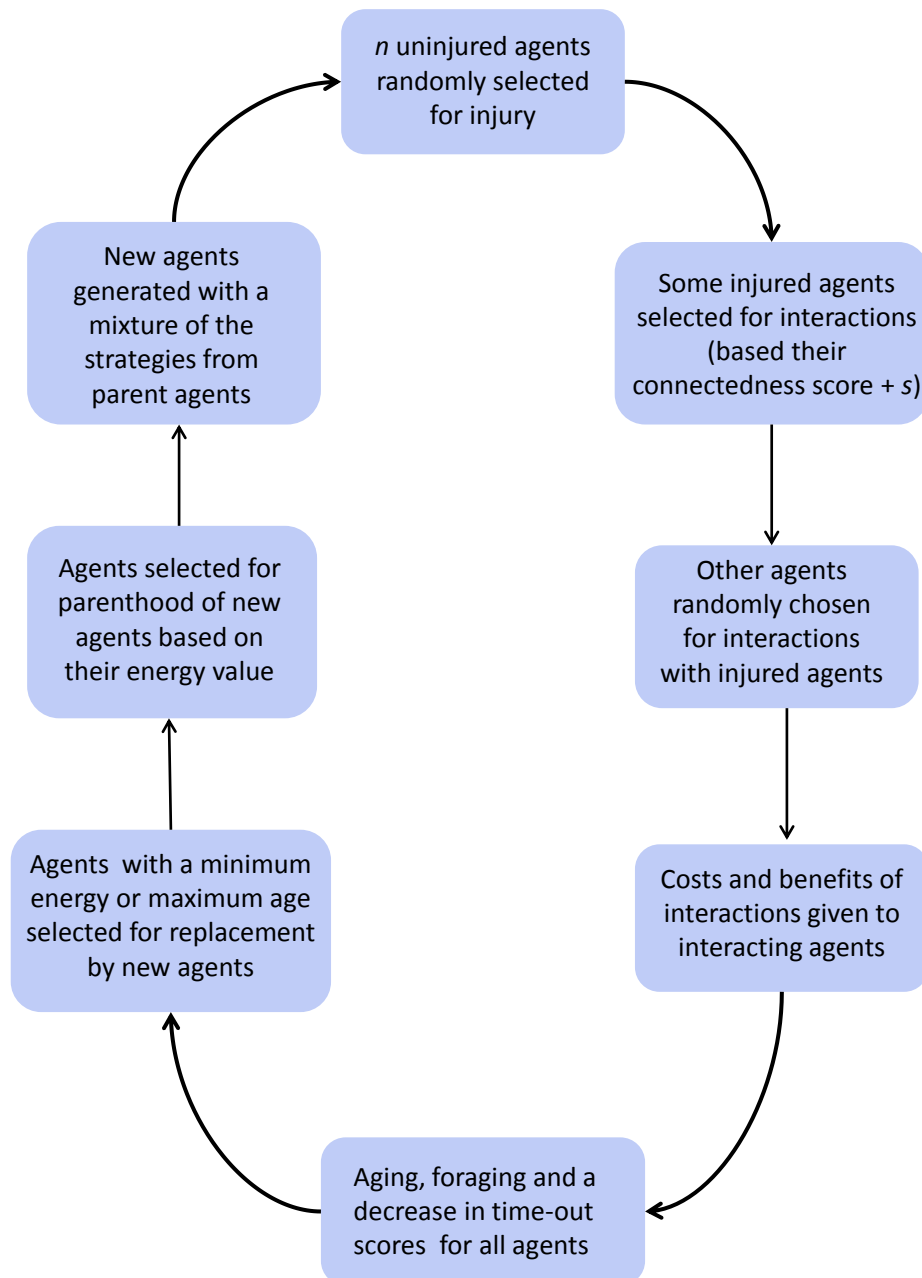


Figure 5: A diagram of processes involved in the model

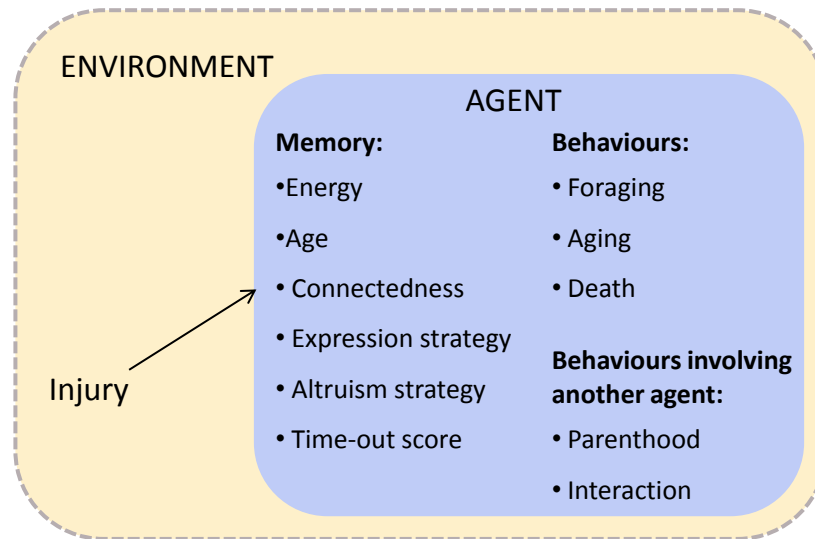


Figure 6: The memory and behaviours of an agent, with environmental influences also shown.

and receive an energy increase. An agent’s energy will increase until it has reached the maximum energy level.

The idea of energy in the model is related to the fitness of an agent; a high energy level means being chosen for parenthood is relatively likely, and vice versa if it is low, if the energy level is zero then parenthood is impossible. Foraging is at a regular rate for simplification reasons, and the energy intake from this increases fitness, as it would in nature.

The age limit means that new generations of agents are made relatively regularly, this is especially important when considering parameters for which energy scores very rarely reach zero.

Connectedness

When an agent is generated it will take a “connectedness” value from an exponential distribution with a mean of 1, these values will be normalised between minimum and maximum values. As such, there will be few agents with a high connectedness value, and

most with a low value.

The connectedness values correspond to the degree of sociability agents have, and the distribution of these values relate to that in a social structure. Injured agents will be chosen to interact based on their connectedness score, for example if an injured agent's connectedness score is 0.7 it will have a 70% chance of being chosen. This chance can also be shifted to model a more sociable population; this will be explained in Section 4.4. This relates to the idea that individuals with more social connections are more likely to be aided when injured. Hence, minimum and maximum connectedness values were necessary to impose as otherwise some agents would always interact when injured and for others it could be impossible.

If an agent helps an injured agent then its connectedness score will be increased and if it ignores an injured agent the score will be decreased, this relates to the idea of an individual's reputation.

Strategies

An agent will have two predefined strategies which determine how it will act in interactions with other agents; one strategy will dictate whether it will express pain when injured, and the other strategy will define whether it will help an injured agent if it is selected to interact with it. Thus, the four different combinations of strategies that an agent can have are:

- Expressor and altruistic,
- expressor and selfish,
- non-expressor and altruistic,
- non-expressor and selfish.

The initial population will have an equal proportion of each of these strategies, i.e. 25% of each.

Hence, if an altruistic agent helps an injured agent then its connectedness score will increase, and thus if it is then in pain itself it will have an increased chance of getting into

a possibly helpful interaction. Similarly, if a selfish agent interacts with an expressor then it will have a decreased chance of getting help if it later becomes injured. Consequently, the connectedness score could also be seen as the reputation of an agent, where good deeds are recorded and rewarded.

Time-out

Unaided, an agent in pain will not be able to forage for a predefined number of iterations (corresponding to a recuperation period), after these iterations the agent will be defined as healthy again. However, if an injured agent is helped then its time-out from foraging will be set back to zero, corresponding to social support aiding recovery.

The initial agent population will all be healthy, hence the “time out” scores will all be zero.

3.2 Pain and interactions

Each iteration a number of healthy agents will be randomly selected to be injured and put into states of pain. It is assumed that all individuals in a population are equally likely to be injured. Out of all the agents who are in pain for an iteration (there will be some from previous iterations who have not been helped yet), some will be selected to interact with another randomly selected agent. This selection is weighted on their connectedness score, as described in Section 3.1. These injured agents and their randomly selected pair will then interact. The pay-offs for these interactions will depend upon the expression strategy of the injured agent and the altruism strategy of the other interacting agent. These can affect the energy level, the time left in pain or the connectedness score of an agent, and are shown in Table 1.

It is worth pointing out that an interaction between an injured non-expressor and any other agent has a zero cost and benefit for both parties, and always results in the injured agent waiting the full time-out from foraging.

To illustrate the cost and benefit effects of interactions, Figure 7 shows components of the lifetimes of four different agents. Each of these agents face interactions at some point, and the effects of these can be seen in their energy, time-out or connectedness levels.

		An injured agent	
		Expressor	Non-expressor
The other agent	Altruistic	An energy cost of expressing (c_{exp})	0
		“Time-out” decreased	
		An energy cost of helping (c_{alt})	0
		A connectedness increase (b_{alt})	
Selfish		An energy cost of expressing (c_{exp})	0
		A connectedness cost (c_{self})	0

Table 1: Pay-offs for the four different interaction possibilities. The injured agent and its pay-offs are shown in boldface and the other interacting agent’s pay-offs are shown in normal face.

3.3 Replacement

Agents will be replaced when they reach the maximum age or have a zero energy score. The new replacement agents will have strategies based on two “parent” agents. As mentioned earlier in Section 3.1, energy level is related to the fitness of an agent, therefore, parents will be chosen based on how high their energy is. New agents will inherit the altruism and expression strategies from their parents (if the parents have different strategies from one another, strategies will be selected randomly). The initial energy level, an age of zero, a connectedness score taken from an exponential distribution with a mean of 1 (and between minimum and maximum connectedness values) and a time-out score of zero will be given to the new agent.

3.4 Parameters

The default parameters of the model can be seen in Table 2. It is not feasible in this study to investigate all of the parameters, thus only a selected few key parameters will be considered. Since this report aims to investigate the influence of social conditions on the success of expression strategies, the key parameters were chosen on the assumption that they reflect social conditions most of all.

Parameter	Default value	Range varied
Number of agents	100	-
Minimum connectedness	0.1	-
Maximum connectedness	0.9	-
Initial energy value	10	-
Maximum energy value	20	-
Energy increment (foraging) each iteration	1	-
Age increment each iteration	1	-
Maximum age	100	-
Number of agents injured each iteration (n)	1	1 - 10
Sociability shift (s)	0	0 - 1
Energy cost of expressing pain (c_{exp})	1	0 - 20
Energy cost of altruism (c_{alt})	1	0 - 20
Connectedness cost of being selfish (c_{self})	0.5	0 - 1
Connectedness benefit of being altruistic (b_{alt})	0.5	0 - 1
Time-out from injury (t)	50	0 - 100

Table 2: Default parameter values. The key parameters chosen to be varied are shown in bold and the range in which they are varied is given.

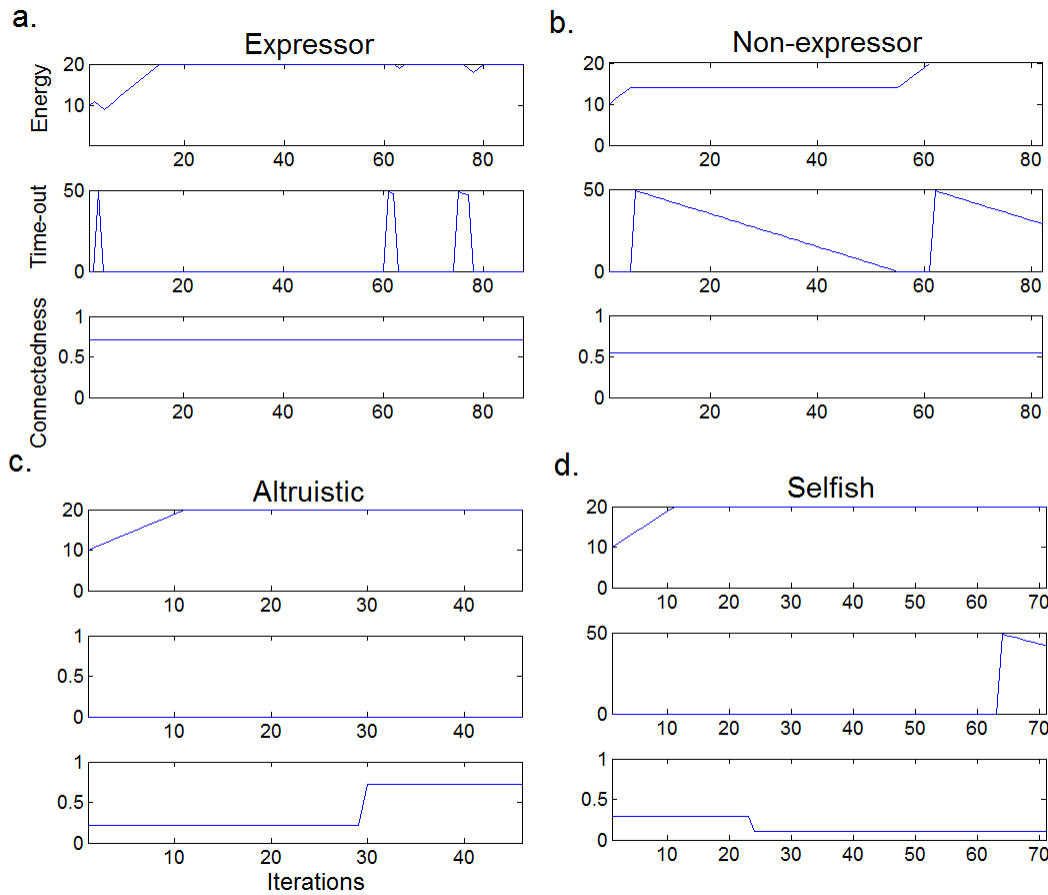


Figure 7: a. An example of an expressing agents’ lifespan. The agent is injured three times, each resulting in an energy decrease and a time-out increase. The first time it is helped immediately, but later it is in pain for a few iterations before help. b. An example of a non-expressing agents’ lifespan. The agent is injured, resulting in a time in pain increase and the agent missing out on 50 iterations of foraging, and then soon after recovery the agent is injured again. c. An example of an altruistic agents’ lifespan. The agent is chosen to interact with an expressing agent in pain, which results in a connectedness increase. d. An example of a selfish agents’ lifespan. The agent interacts with another in pain, which results in a connectedness decrease as the agent does not help. Later the agent is injured. The agents shown here have started from different ages and all died from “old age”.

In this report the cost of expression (c_{exp}), the cost of altruism (c_{alt}), the cost of selfishness (c_{self}), the benefit of altruism (b_{alt}) and the time in pain (t) will be referred to as “the cost and benefit parameters”.

With a lack of literature to calibrate the model parameters from, a number of assumptions had to be made to find a default parameter set. Firstly, it was assumed that energy costs are equal to the foraging energy increment (i.e. $c_{exp} = c_{alt} = 1$). Also, it will be assumed that an agent’s connectedness score is the same as its chance of interacting when injured (i.e. $s = 0$) and that 1% of agents will be injured each iteration ($n = 1$).

The other cost and benefit parameters were taken as their mid-value. This was justified by a preliminary analysis of parameter sensitivity in which the model was ran with low, mid and high values for t , c_{self} and b_{alt} . The results of this can be seen in Figure 8. From these graphs it can be seen that c_{self} and b_{alt} seem to make little impact on the strategy outcomes, hence a mid value may as well be used. Increasing t has a more obviously effect on the results; when t is high there are more expressing altruists. This increase appears to be linear and in Figure 9 it is shown that using a mid-value for t gives a reasonable number of agents in pain in each iteration. Hence, a mid-value for t will be taken.

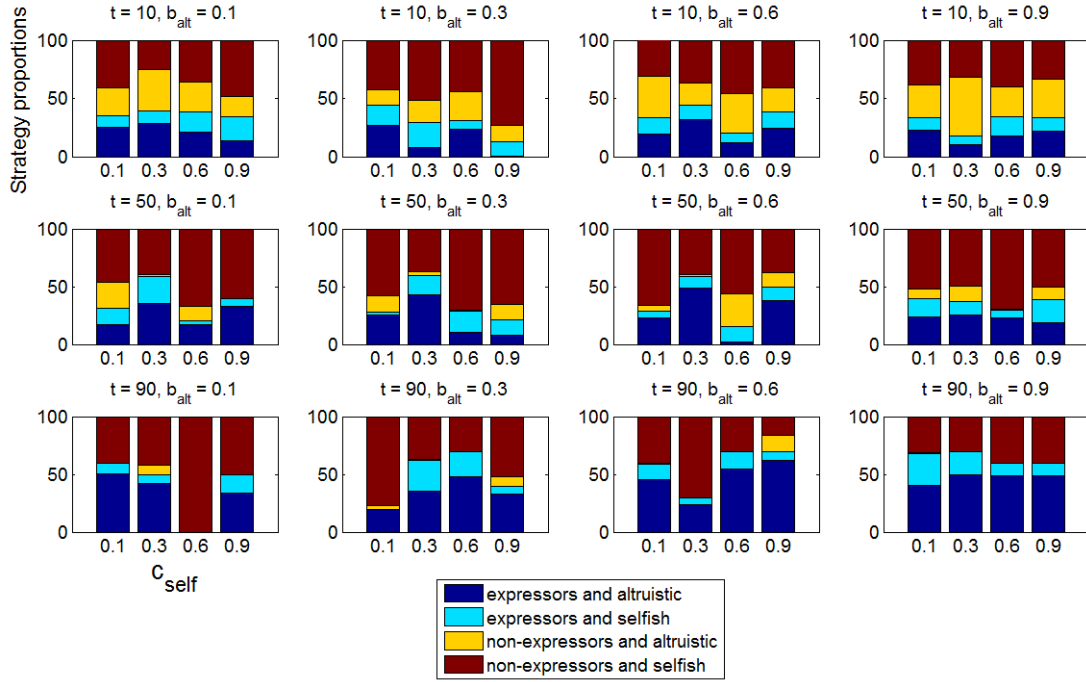


Figure 8: The proportions of strategies in the population at each iteration where the time in pain (t), the connectedness benefit of altruism (b_{alt}), and the connectedness cost of selfishness (c_{self}) are varied between low, mid and high values. Found using $c_{exp} = c_{alt} = 1$ and ran for 3000 iterations averaged over 10 trials.

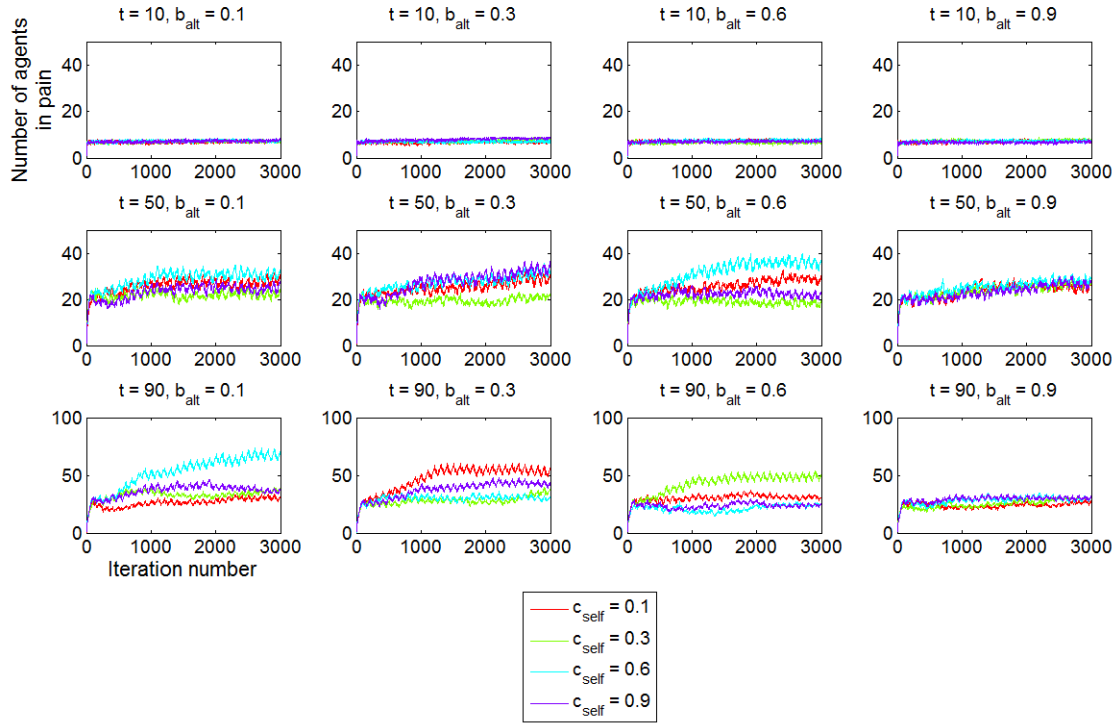


Figure 9: The number of agents in pain when varying the time in pain (t), the connectedness benefit of altruism (b_{alt}) and the connectedness cost of selfishness (c_{self}). The cost of expression and cost of altruism were both set as 1. Found using 3000 iterations and 10 trials.

4 Experiments and results

By using the model to look at the proportions of strategies over many iterations, and how these change in different circumstances, the conditions for which pain expression is favourable can be predicted.

Hence, to test these conditions the model was run with a variety of parameter values (these can be seen in Table 2). However, before this could be done, preliminary experiments were undertaken. These included testing the model’s sensitivity to the initial agent population, finding how long to run each of the simulations for and finding how many repeated simulations (which will also be referred to as “trials”) are needed to be averaged over in order to get consistent results.

After this preliminary analysis, the effect of the cost and benefit parameters were looked at, and then parameters which corresponded to more direct environmental changes were varied. Some basic results of these changes are described in this section, however the majority of the analysis is done in Section 5.

4.1 Preliminary experiments

4.1.1 The initial agent population

Due to the time constraints of this report it is preferable to only consider the outcomes of one initial agent population. However, this can only be justified if the initial agent population chosen makes little difference to the results. This was tested by looking at the outcomes of both 10 repeated simulations of a single, default initial agent population and the outcomes of 10 different initial populations. The default parameters from Table 2 were used and results averaged over 100 repeated simulations ran for 10,000 iterations.

The results are shown in Figure 10. A two-sample t-test resulted in the acceptance of the null hypothesis that the data from both these sets of results have equal means and variances. Hence, it is likely that most of the variance seen is due to the stochasticity of the model, rather than the stochasticity in finding the initial agents. Therefore, the results found in this report will use the same, randomly selected, initial agent population in every simulation. This will be referred to as the “default initial agent population” from



Figure 10: The final proportions of strategies in the population run 10 times using the same default initial agent population (left) and run using 10 other initial populations (right). Found using the default parameters from Table 2, 10,000 iterations and 100 trials.

this point.

4.1.2 Simulation variability

It is necessary to test how much of an effect the stochastic elements of the model have on the variability of the results. This was done by looking at the outcomes of running the model 100 times over 10,000 iterations. The default parameters from Table 2 and the default initial agent population were used.

Figure 11 shows a section of these results, the rest can be seen in Figure 21 in Appendix A. Results show that 24% of simulations ended in altruistic expressor strategies, 4% as selfish expressors, 6% as altruistic non-expressors and 60% as selfish non-expressors. Furthermore, 4% finished in a mixture of non-expression strategies and the remaining 2% in a mixture of the expressor strategies. Hence, results vary significantly and therefore simulations need to be left running for long enough for steady behaviour to develop and furthermore each simulation needs to be rerun enough times for the average result to describe the behaviour well.

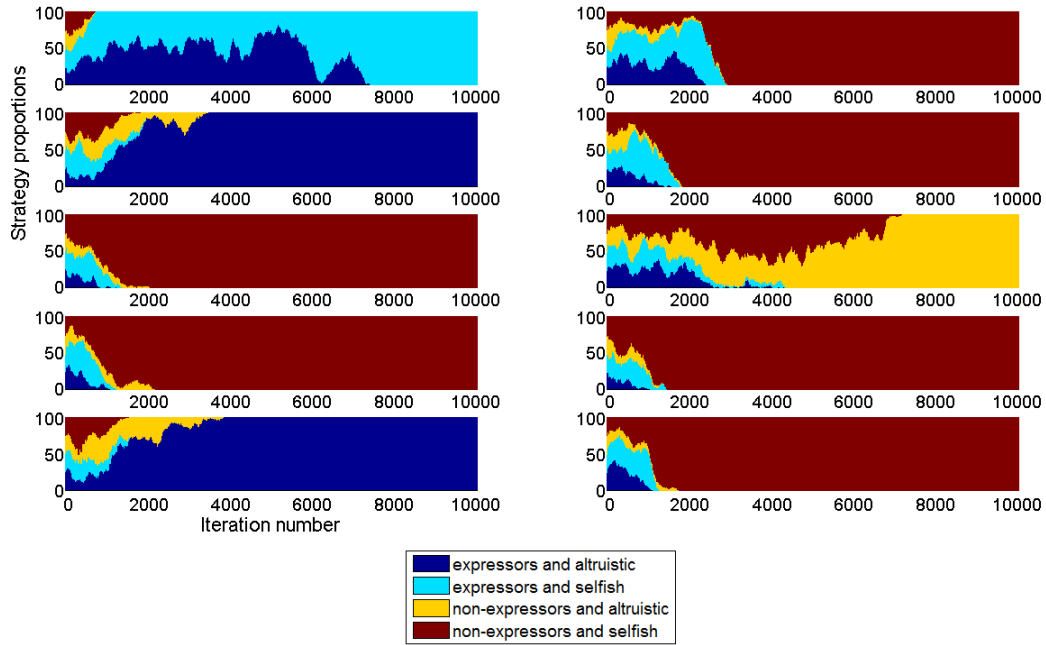


Figure 11: The results of ten simulations over 10,000 iterations, starting with the same initial conditions and using the default parameters from Table 2.

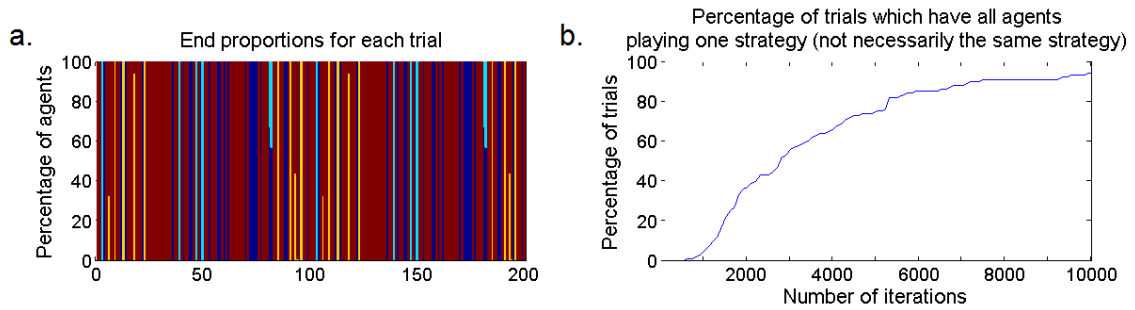


Figure 12: a. The proportions of strategies in the last iteration from 200 trials. b. The percentage of repeated simulations which end in all agents playing the same strategy after leaving the model for different numbers of iterations. Found using the default parameters from Table 2.

4.1.3 The number of iterations

In order to make sure results are taken from simulations which show steady behaviour, the simulations need to be left running until most of the trials have reached a one-strategy outcome. When the whole agent population only uses one strategy then no other strategy can be used again. To test how long simulations need to be left running for, the percentage of trials which finish in a one-strategy outcome were found at each iteration using the default parameters from Table 2.

Figure 12a shows the different strategies used in the 10,000th iteration for 200 repeated simulations, it can be seen that most of these repeats have finished in a one strategy outcome. Figure 12b shows how leaving the simulation running for longer means that end results are more likely to have all agents using only one strategy. After 10,000 iterations it can be seen that almost all of the repeated trials will end in a one-strategy outcome. Hence, the end strategy proportions when trials have run for 10,000 iterations will be used.

4.1.4 The number of repeated simulations

The other factor necessary to consider is the number of repeated simulations needed to be ran in order to get realistic average strategy proportions. Although finding the average of more simulations is more accurate, it is time consuming, so it is useful to find an equilibrium point. This was tested by finding the proportion of strategies at the 10,000th iteration by averaging over different numbers of trials using the default parameters from Table 2.

Figure 13 shows the final strategy proportion outcomes found by averaging up to 200 trials. It can be seen that from an average of about 100 repetitions the proportions stay reasonably similar to the proportions seen when the number of repetitions is increased.




Figure 13: Strategy proportions found when averaging over different numbers of repeated simulations. Found using the default parameter set from Table 2 and run for 10,000 iterations.

Figure 14 shows that after 10,000 iterations and by averaging over 100 trials there is a steady behaviour in the proportions of strategies used. Thus, simulations will be ran using these values.




Figure 14: Strategy proportions at each iteration, averaged over 100 trials. Found using the default parameter set from Table 2 and run for 10,000 iterations.

4.2 Cost and benefit effects

The cost and benefits to pain expression and altruism can vary for different species, or in different environments. Thus, by varying them in this model, their effect on the strategy proportions outcomes can be understood.

To investigate this, the cost and benefit parameters (c_{exp} , c_{alt} , c_{self} , b_{alt} and t) were

varied 21 times each (these parameters and the ranges varied can be seen in Table 2). One parameter at a time was varied and the others were kept as their default values. For each variation the proportion of each of the strategies in the 10,000th iteration was found, averaging over 100 trials. Unless otherwise stated default parameters were used.

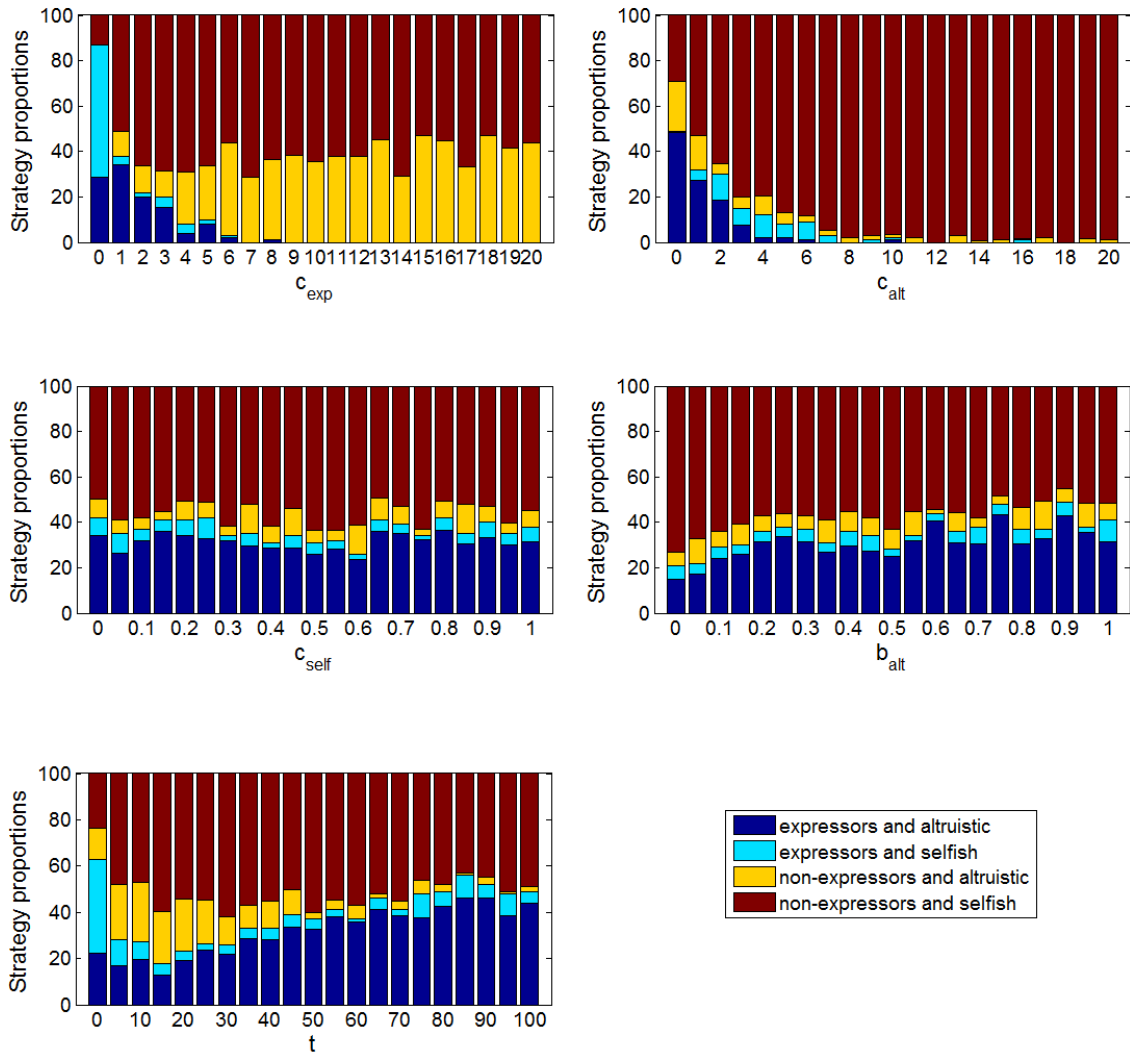


Figure 15: The effect of changing parameters on the final strategy proportions. Results were found from an average of 100 trials at the end of 10,000 iterations, and unless otherwise stated parameters are taken from the default set shown in Table 2.

c_{exp} and c_{alt} were both varied between having no cost and having a cost of 20. If one of these parameters was greater than or equal to 20 then any agent which expresses pain or is altruistic would die, thus investigating the effect of this cost at values higher than 20 is unnecessary as the results would be the same. Similarly t is varied between 0 and 100 because if the time in pain is greater or equal to 100 then any non-expressing agent in pain would always die before recovery.

The results from these variations can be seen in Figure 15. Appendices B to F show the average proportions at each iteration and the proportions averaged over different numbers of trials for each of the variations. From these it can be seen that they have all reached reasonably steady behaviour, thus supporting the results found in the preliminary experiments.

The graphs for c_{exp} and t show quite logical and simple results; when the cost of expression increases there is less expression and when the cost of being a non-expressor is higher (i.e. the time in pain is higher) there is more expression. However, when looking closer there are other, more complex, forces acting on the strategy proportions, these will be described more in Section 5.2.

Looking at the results for varying c_{alt} there is not just an increase in the cost of altruism leading to less altruism. Instead there is also an indirect effect on the number of expressors, where even expressing selfish agents decrease in number. This is because if there are less altruistic agents because of a high cost, then being an expressor is no longer as beneficial as it is more unlikely to be helped.

Varying c_{self} makes little difference to the outcomes, and varying b_{alt} results in a very small increase in expressing altruists as b_{alt} is increased. This could be because connectedness does not have very much influence in the model.

4.3 Danger

In a more dangerous environment it would be reasonable to assume that more individuals would be injured. Hence, to test the effect of a dangerous environment on the proportions of agents predicted by the model, the number of agents injured each iteration (n) was varied between 0 and 10. The proportions of strategies in the 10,000th iteration were found by averaging over 100 trials. Other than n , the default parameters from Table 2 were used.

The results of these changes are shown in Figure 16. Figure 32 in Appendix G shows the number of agents in pain and the number of agents in pain who are interacting at each iteration for each of the changes. Appendix G also shows that results are reasonably steady when averaging over 100 trials and running simulations for 10,000 iterations.

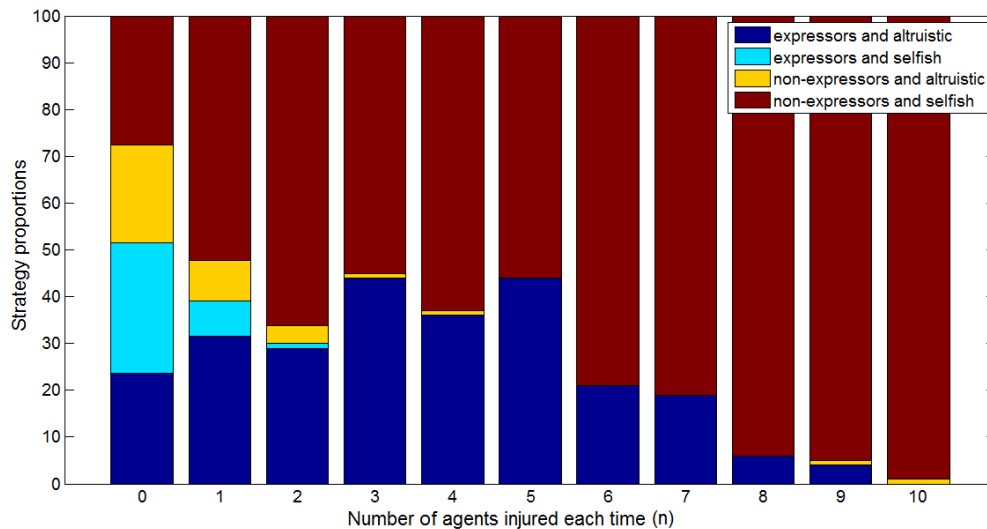


Figure 16: The effect of changing the number of agents injured each iteration (n) on the final strategy proportions. 10,000 iterations ran and 100 trials averaged, other than n parameters were taken from the default set shown in Table 2.

Increasing n shows that there is a general decrease in the number of simulations ending with a population of expressors. This corresponds to the decrease in the benefit of expression because there is more chance of getting injured again. Appendix G shows that

when n is high (> 8) almost all of the population is injured.

4.4 Sociability

Assuming that in a more sociable community an individual will have more chance of being observed by another individual if it is in pain, then sociability should have an effect on the benefits of pain expression as it would increase the chance of, possibly helpful, interactions. To test this effect on the outcomes of the model, the chance of interaction will be increased.

In the previous analysis an agent in pain was chosen to interact with another directly based on its connectedness score (it was mentioned in Section 3.1 that if an agent had a 0.6 connectedness score it has 60% chance of interacting). To test the effect of sociability, this score will be shifted by s . Hence, agents will now all have an increased probability of their connectedness score $+ s$ of interacting. The proportions of strategies in the 10,000th iteration were found by averaging over 100 trials. s was varied between 0 and 1. Other than s , the default parameters from Table 2 were used.

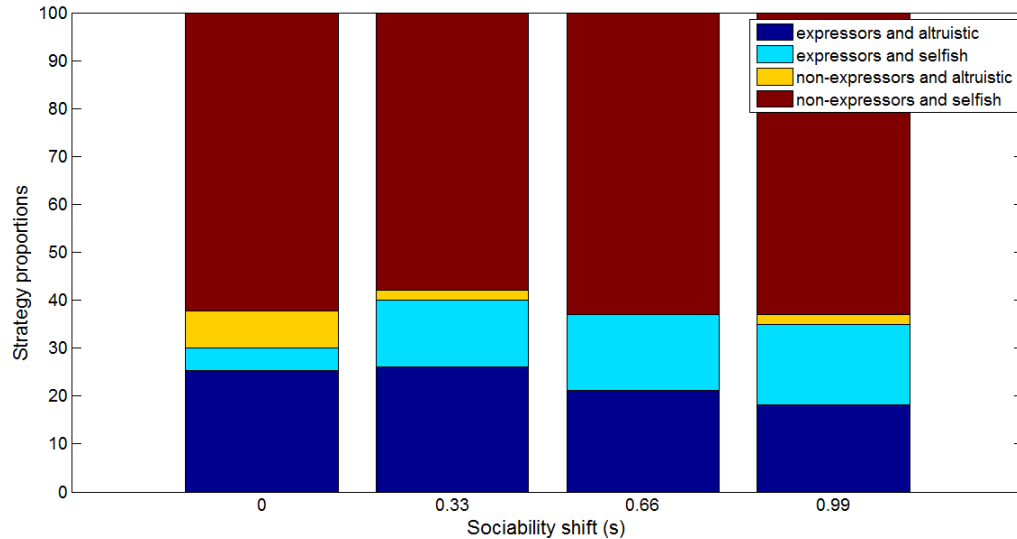


Figure 17: Varying the shift in sociability, s . 10,000 iterations ran and 100 trials averaged, parameters were taken from the default set found in Table 2.

The results of these sociability changes are shown in Figure 17. The graphs in Appendix H show that results are reasonably steady when averaging over 100 trials and running simulations for 10,000 iterations.

The results show that when increasing s there is a small increase in the number of selfish expressors and a decrease in the number of altruistic expressors in the final proportions. This could be because the relative benefit of altruism is decreased as all agents have more chance of interacting regardless of strategies.

4.5 Antagonists

Another cost to expressing pain is its attention to vulnerability in the presence of antagonists. To incorporate this idea into the model, selfish agents were also able to steal energy from injured agents in an interaction. During this type of interaction the expressing agent will lose a certain amount of energy which the selfish agent will gain. This idea could correspond to a predator attack or an instance of resource stealing, such as kleptoparasitism. When the energy stolen is equal to the expressing agent's energy level then it will be left with no energy and will die. The proportions of strategies in the 10,000th iteration were found by averaging over 100 trials using the default parameters from Table 2.

The effect of varying the amount of energy stolen is shown in Figure 18. The graphs in Appendix I show that results are reasonably steady when averaging over 100 trials and running simulations for 10,000 iterations.

The graphs show a clear effect of including energy stealing between an expressing agent and a selfish agent; when more energy is stolen there is less expression. This is an obvious result as there is an increased cost to expression.

4.6 Inheriting connectedness

The results thus far have shown that the connectedness values are not very influential on the outcomes of the model. To try to change this, the assumption that connectedness is something randomly assigned to an agent was lifted, and instead connectedness was inherited. In this situation offspring would have an average of its parents' connectedness values. Using the anti-Lamarckism idea that a characteristic which has changed in an




Figure 18: Varying the energy stolen in an expressor/selfish interaction. 10,000 iterations ran and 100 trials averaged, other parameters are taken from the default set shown in Table 2.

individual’s lifetime should not be inherited, c_{self} and b_{alt} were set to zero. To test what difference this would make, the model was ran with the default parameters from Table 2. Furthermore, a small range of the values for c_{exp} , c_{alt} and t were varied.

Figure 19 shows the results found when connectedness can be inherited. The graphs in Appendix J show that results are reasonably steady when averaging over 100 trials and running simulations for 10,000 iterations.

The results show that the outcomes when connectedness is inherited are very similar to those for when it is not. However, there does appear to be a slight increase in the amount of selfish expressors and a decrease in altruistic expressors when connectedness is inherited.

It might be expected that if $c_{alt} = 1$ (i.e. the energy cost of altruism is low) then there would be more altruistic agents than selfish agents in the population. However, there is an equal number of these strategies when connectedness is inherited. Hence, perhaps inheriting connectedness gives less reliable results than when it is not inherited.




Figure 19: Strategy proportions found for each iteration when connectedness can be inherited (a) and the last iterations when connectedness is not inherited (b). Found using the default parameter set from Table 2 and averaged over 100 trials.

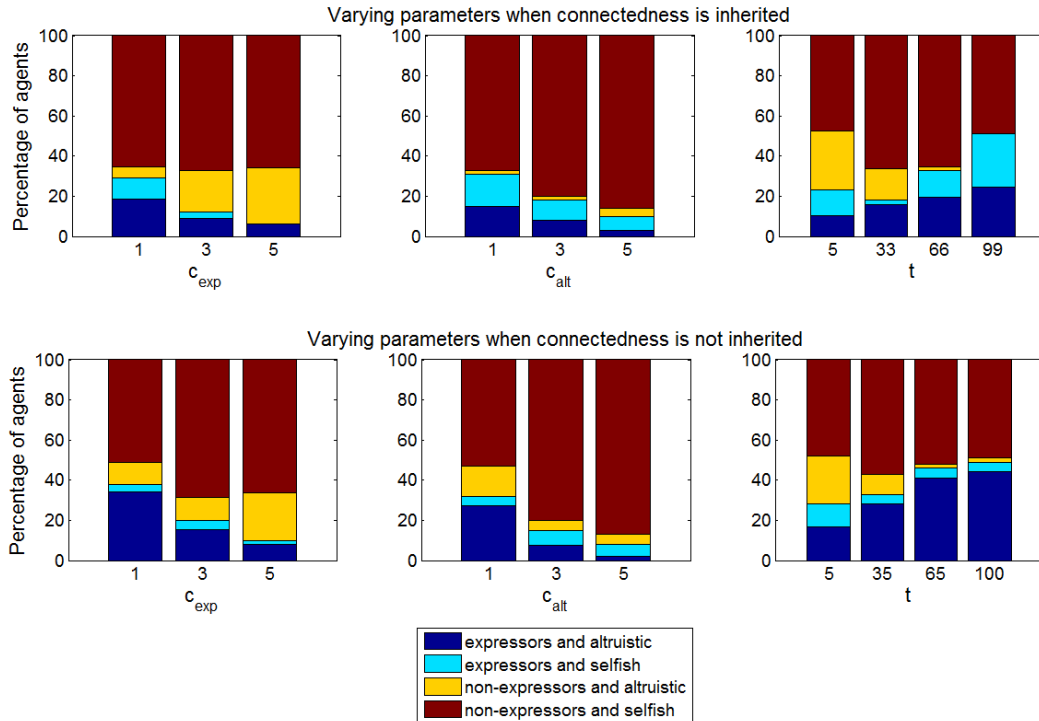


Figure 20: Varying some parameters in a model where connectedness is inherited and $c_{self} = b_{alt} = 0$ (above). For reference, the results for when connectedness is not inherited and $c_{self} = b_{alt} = 1$ are shown below. 10,000 iterations ran and 100 trials averaged, unless otherwise stated, parameters are taken from the default set found in Table 2.

5 Discussion

It is important to remember when interpreting the results that the proportions shown on the graphs are averages of the end proportions of different simulations (for example, see Appendix A for the individual proportions of trials from the default parameters). Furthermore, most of these end points are just of one strategy. So, for example, in the graph where c_{exp} is varied in Figure 15 the proportions at $c_{exp} = 20$ do not show that the final population had roughly 50% non-expressing altruists and 50% non-expressing selfish agent, it means that roughly 50% of repeated simulations ended in mostly all non-expressing altruists and 50% in non-expressing selfish agents.

It is also important to remember the stochasticity seen in the results, and therefore not to give too much weighting to small changes as these can arise through the model's natural variance, rather than an actual effect of a parameter change.

This section will be divided into some general observations about the results of the model and then a more in-depth analysis of the results of each experiment.

5.1 General observations

It was found of the expressing agents that altruists tended to do better than selfish agents and of the non-expressors selfish agents did better in most of the conditions looked at. This result reflects the coevolution of pain expression and its ability to be recognised (and therefore help given). Hence, the model predicts that strategies of expression will exist in situations of social support. Thus, the result is compatible with the evolutionary hypothesis that in the presence of potential caregivers there will be more pain expression (Williams, 2002).

For the success of selfish expressors, there is a kind of stabilising feedback loop operating. If their numbers increase they become less likely to interact with an altruistic agent, and hence expressing becomes less beneficial so their numbers will decrease, but this in turn makes their strategy more beneficial so their numbers can increase, and so on. Hence, there is less chance of selfish expressors dominating the population than there is for altruistic expressors.

Altruistic non-expressors tend to only do well if there are less expressors, as otherwise they pay a cost and never receive a benefit. In the absence of expressors altruistic and selfish non-expressors pay no cost and receive no benefit, hence their proportions will stay at more or less the same level even when the environment changes.

In the absence of expressors, the number of selfish and altruistic non-expressors in the last iteration will be determined by a process similar to genetic drift, i.e. random fluctuations will determine the proportions at the 10,000th interval. In terms of an average over many trials, it can be seen (see Appendices B to J) that the proportions of these strategies remain the same as they did when expressors were still in the populations. This may be because the chance of a simulation ending in one of the non-expression strategies by random fluctuations would be increased if the majority of agents had this strategy when expressors were in the population. For example, if a population had 25% altruistic non-expressors and 75% selfish non-expressors immediately after expressing agents died out, it may be expected that selfish non-expressors would be 3 times more likely to dominate the population than altruistic non-expressors. This may explain why even when there are no expressors the population the proportion of selfish non-expressors is higher than the altruistic non-expressors.

The model also indicates that in situations of an increased danger (for example, increasing the number of agents injured each iteration or included an antagonist agent), there would be a decrease in the number of expression strategies. This fits with the effects of social threat discussed in Section 2.1.2.

5.2 Cost and benefit parameters

The results for varying c_{exp} show a decrease in the number of expressors when the cost of expressing is higher.

Further to this, when c_{exp} increases there are increasingly more non-expressing altruists until the point when c_{exp} is too high for expressors to exist at all (when $c_{exp} = 8$). This may be due to the fact that a non-expressing altruist will pay the cost of helping others, but will never receive any benefit back (as it never signals that it is in pain). Hence, the presence of expressors in a population is detrimental to a non-expressing altruist, so when expressor

numbers dwindle the non-expressing altruist can thrive. When there are no longer any expressing agents left in the population there will never be any interactions. Hence, selfish or altruistic non-expressor strategies will occur in roughly the same proportions as they did in the iterations when expressors were around (as explained in Section 5.1). Furthermore, the effect of increasing c_{exp} obviously makes no difference when there are only non-expressors around.

When $c_{exp} = 0$ practically all of the repeated simulations end in an expressor strategy. This is reassuring as obviously being a non-expressor when $c_{exp} = 0$ would be less advantageous. However, there are also roughly equal numbers of both altruistic and selfish expressors. This could be because the costs and benefits of altruism and selfishness are balanced, and since there is no cost to expression we see that it is equally likely for a simulation to end in either of these strategies. Expression becomes less influential and non-expressors more so when c_{exp} is increased to 1. Although the number of altruistic expressors stays at a similar level, the number of selfish expressors decreases significantly, this could be because of the feedback loop mentioned in Section 5.1.

Varying c_{alt} shows a decrease in the number of altruists in the final proportions when this cost is increased. There is also a decrease in the number of selfish expressors. This can be explained as an indirect effect where by increasing c_{alt} there are less altruists and therefore less expressors as expressing is not as beneficial in a population with few altruists. Small proportions of altruistic non-expressors are seen even when c_{alt} is very high. This is because (as it was in the case of a high c_{exp} value) if there are no expressors then an agent's altruism strategy is irrelevant.

The results show that the effect of connectedness in the model does not seem to make much of an impact. It can be seen by varying c_{self} and c_{alt} that the strategy proportion outcomes are practically unchanged.

Increasing t results in an increase in the proportion of expressors in the final iterations, which is because t is a cost to not expressing, therefore when it is high there will be less non-expressors. Increasing t also reveals a similar, but reversed, effect to when varying c_{exp} ; when expressing is less beneficial there are less expressors and therefore more altruistic non-expressors. However, the time in pain can never be low enough to result in only

non-expressor strategies, so there will always be some expressors.

When $t = 0$ it is found that the end proportions are relatively equal, this is because there are effectively no injuries and therefore no interactions. An agent who is injured will have a time-out score of zero, and therefore it will be classed as healthy. Thus the end strategies will arise through a process similar to genetic drift.

5.3 Danger

When varying the number of agents injured each iteration the number of expressors roughly decreased when n was increased. This could be because an agent could expend energy expressing pain, get helped, and then soon after become injured again. Hence, if agents are becoming injured regularly then it might be more beneficial to not express pain as they will be in pain most of the time regardless. This result could be compared to the ideas of stress-induced analgesia; in times of more danger it could be more beneficial to conserve energy and stay alert to further danger rather than to express pain (Langford *et al.*, 2011).

There is also a trend in the number of altruistic expressors; this strategy seems to increase in numbers and then decrease. This may be because for low-to-medium danger environments ($1 \leq n \leq 5$) being altruistic will be beneficial because they will be more likely to face interactions in an increasingly dangerous environment. Then, after a threshold ($n \approx 6$) there are many injured agents in the environment, so altruists are consistently paying helping costs. Hence, another influence on expression becoming less prominent is because of the decrease in altruists who can help them. Perhaps this could be related to an idea of “giving up” where help is given when not too many others are in pain, but when many are in pain an altruistic individual gives up trying to help.

When $n = 0$ there are no injuries and therefore no interactions, hence every strategy is equally likely to dominate the population.

5.4 Sociability

In a more sociable environment the model predicted that there would be an increase in selfish expressors and a decrease in altruistic expressors.

This could be because the chance of interaction is increased for both altruistic and selfish expressors equally and since there is a limit to connectedness for higher values of s most agents have close to the limit for connectedness. Hence, the benefit of altruism becomes less influential as s increases. When $s = 0.99$ the benefit of altruism is effectively zero, as for this every agent has a 0.9 chance of interaction regardless of strategy. However, it is still seen that both altruistic and selfish agents both exist in roughly equal proportions. Perhaps this could be explained by the stabilising feedback loop idea mentioned in Section 5.1, where the success of selfish expressors is restricted.

Counterintuitively, when the sociability is increased the proportion of expressors stays the same and furthermore non-expressors make up the majority of the strategy proportions. Alternatively, it might have been expected that when there is more chance of being helped, expression would be most beneficial, and hence a more drastic increase in the proportion of expressors would have been seen.

5.5 Antagonists

When energy can be stolen there is an extra cost to expressing. Hence, a decrease in the number of expressors is seen when the amount of energy stolen is increased. There is also an extra benefit to selfishness when energy is stolen, but the effects of this are not seen as altruism strategies are irrelevant in the absence of expressors. Hence, both of the non-expressor strategies are seen when the energy stolen is high.

This result corresponds to the idea of expression drawing attention to vulnerability and thus exploitation (Williams, 2002; Wiech & Tracey, 2013).

5.6 Inheriting connectedness

When connectedness is inherited the model predicts an increase in the number of selfish expressors and an increase in the number of altruistic expressors. These differences roughly balance the numbers of altruistic expressors and selfish expressors. However, this result is less intuitive than those found when connectedness is not inherited as for certain changes in parameters it might be expected that there would be a more altruistic expressors than selfish expressors. For example, when the cost of altruism is low, the cost of expression

is low, or the cost of non-expression is high (i.e. t is high), there would be an increased benefit of altruism.

6 Conclusions and extensions

This report aimed to develop a model which can predict whether pain expression is favourable or not in a range of different social environments. This aim has been met and the results correlate with the behaviour seen in some of the experimental research on pain expression, specifically the work on social support by Langford *et al* (2006), and the work on social threat by Langford *et al* (2011) and also by Peeters and Vlaeyen (2011).

The results found in this report confirm logical intuitions about when pain expression would exist; when there is an increase in the physical energy cost of expression there is a reduction of expression and when there is an increase in the cost of not expressing there is an increase in expression.

Other results have picked out the relationship between altruism and expression; namely that when there are few altruists there will also be a reduction in expression. The model also found that in a more dangerous environment there will be less expression. Likewise, in an environment which includes antagonists, a decrease in the number of expressors was found. Hence, these results predict social support motivating expression and an increase in social threat motivating less expression, as discussed by Williams (2002).

The results found in this report have all been analysed visually. Although some of the changes made have shown quite an observable difference, a more in-depth statistical analysis should be made on the results to confirm the significance of the changes seen.

Due to the time limitations of this project not all parameters have been explored. For example, it would have been of interest to look at how the initial proportions of strategies or the energy increment due to foraging affected the number of expressing strategies. In the future, it would also be interesting to extend the model described here by introducing a spatial dimension. This could provide information about the effect of neighbourhoods and boundaries in an environment on the different strategies found. Using these could yield results like the chaotic or clustered configurations of cooperators found when investigating the spatial prisoner's dilemma by Nowak, Bonhoeffer & May (1994). As such, the results could even be used to give insight into the variance of pain behaviour seen in different cultures (Williams, 2002).

Experimental observations about the costs and benefits of pain expression and its detection often point to evolutionary influences as an afterthought. Hence, this report hopes to generate more specific evolutionary-based research questions. These investigations could find exact values for the parameters in this report, and thus lead to explicit, empirical results for the situations for which pain behaviour is expressed. Additionally, and perhaps most importantly, investigations could lead to a reduced application of the operant ideas of pain expression by clinicians, where patients can be undertreated due to scepticism about the amplification of pain behaviour (Williams, 2002) and social influences are disregarded.

This report has developed an agent-based model which predicts when expressing pain becomes beneficial. This approach is entirely novel in an application to pain behaviour and hopes to demonstrate its use and motivate further, more detailed, work and research questions.

Appendix

A. Default parameters




Figure 21: The strategy proportions through 10,000 iterations from each of 100 repeated simulations. Default parameters used.

B. Varying c_{exp}

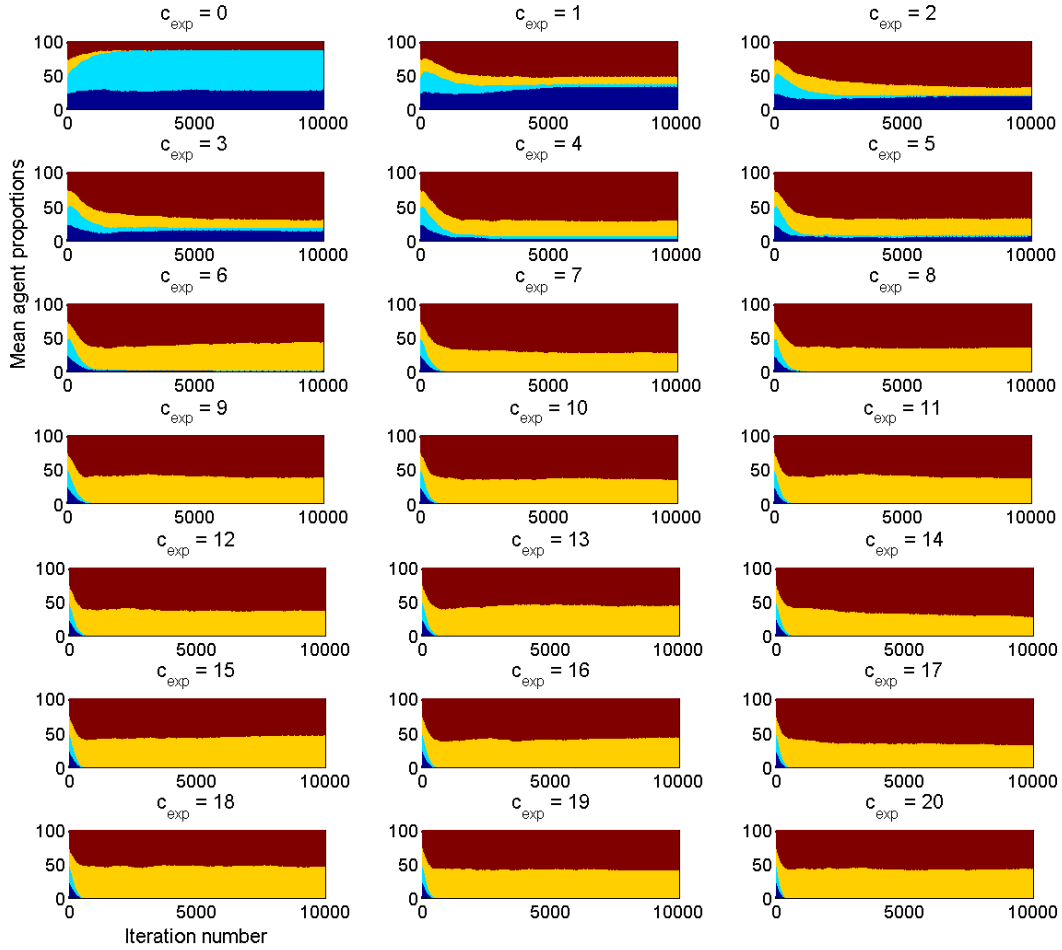


Figure 22: The average strategy proportions at each iteration, varying the cost of expression (c_{exp}). The average of 100 trials at each iteration, default parameters used elsewhere.

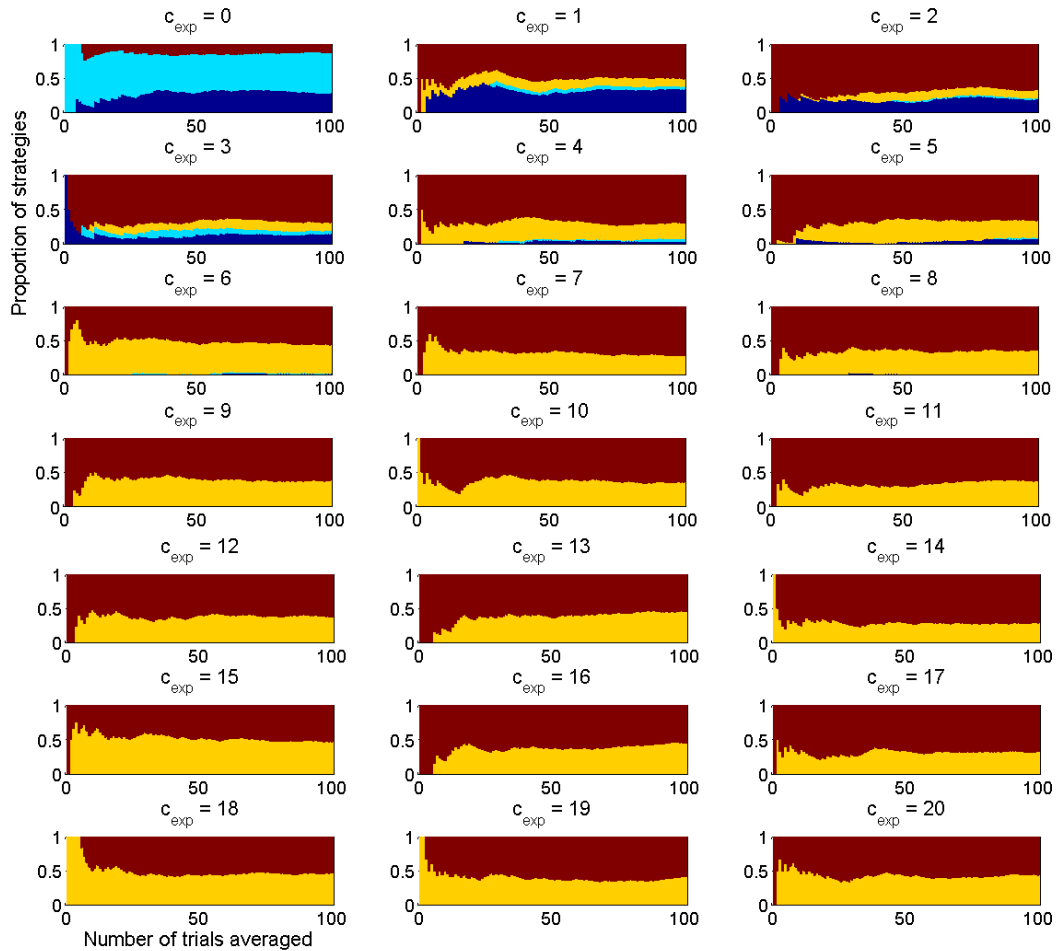


Figure 23: Strategy proportions averaged over different numbers of trials, varying the cost of expression (c_{exp}). The 10,000th iteration averaged over a different numbers of trials, default parameters used elsewhere.

C. Varying c_{alt}

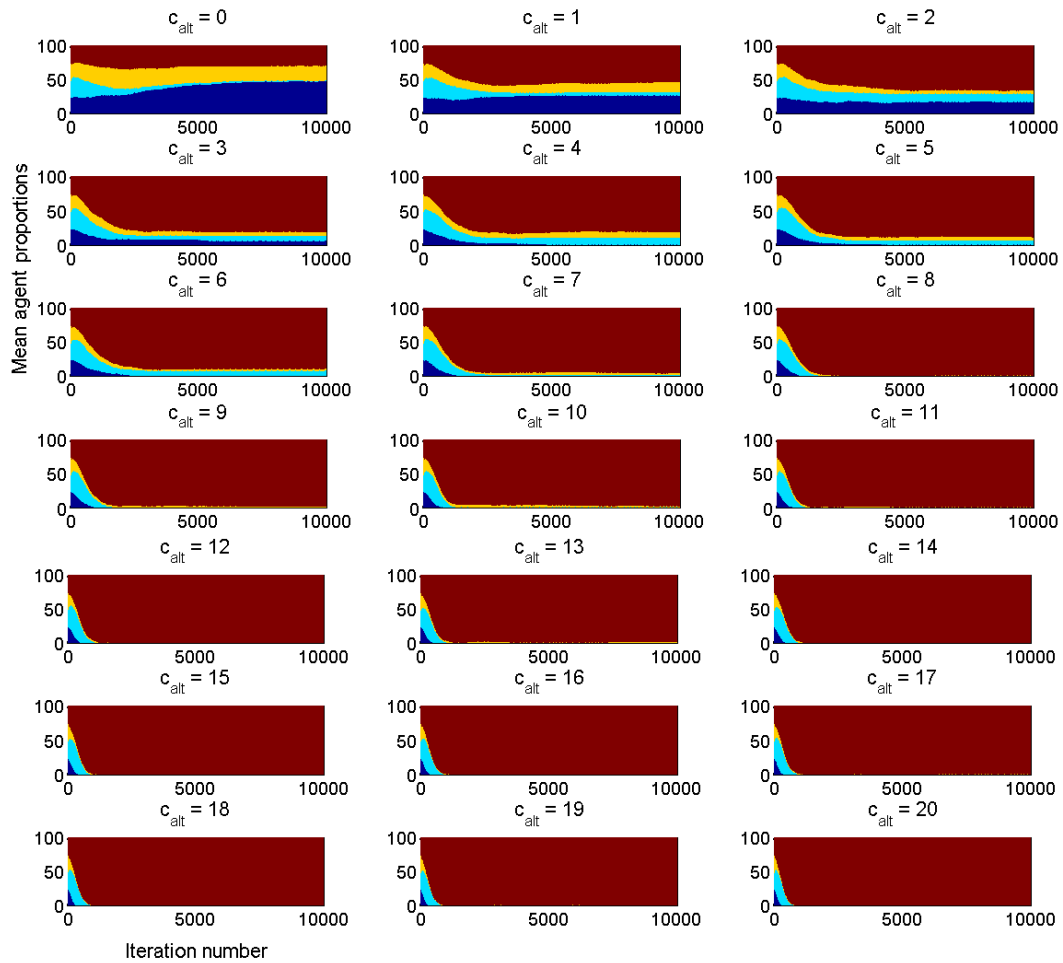


Figure 24: The average strategy proportions at each iteration, varying the cost of altruism (c_{alt}). The average of 100 trials at each iteration, default parameters used elsewhere.

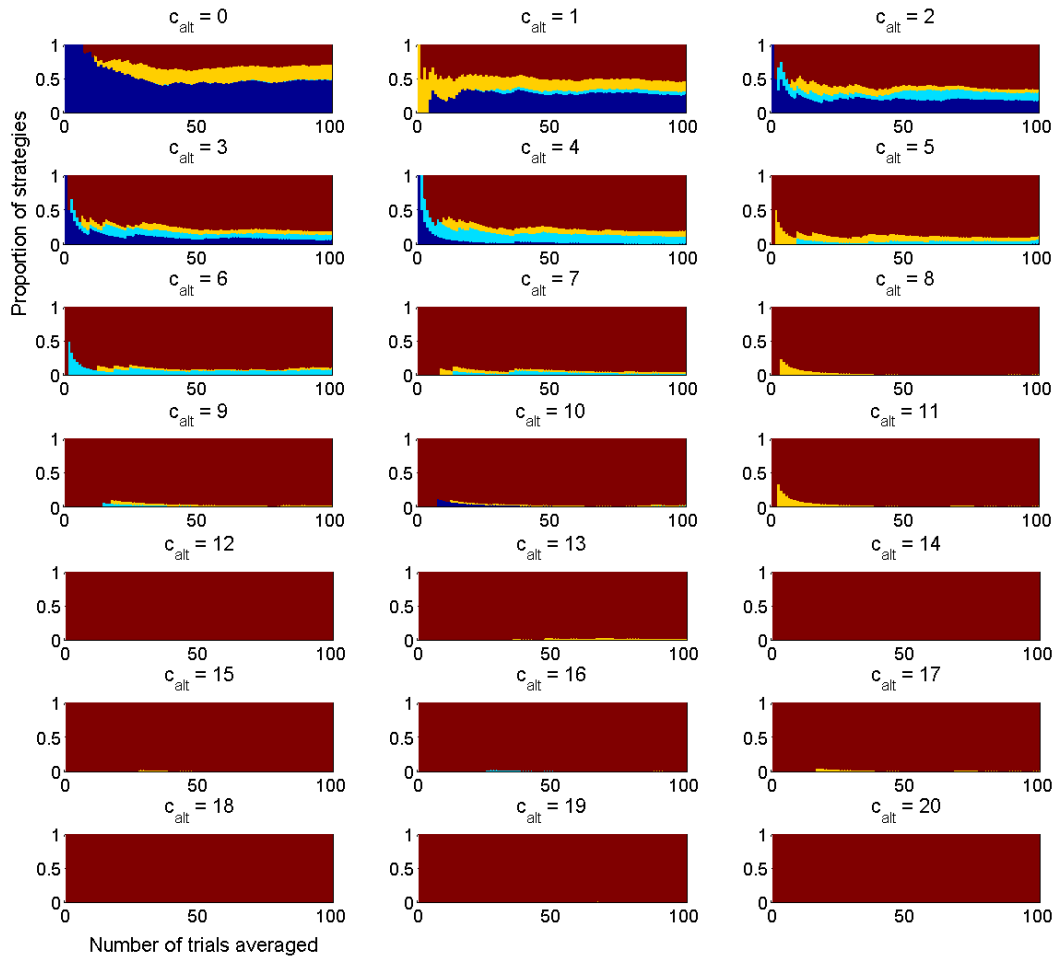


Figure 25: Strategy proportions averaged over different numbers of trials, varying the cost of altruism (c_{alt}). The 10,000th iteration averaged over a different numbers of trials, default parameters used elsewhere.

D. Varying c_{self}




Figure 26: The average strategy proportions at each iteration, varying the cost of selfishness (c_{self}). The average of 100 trials at each iteration, default parameters used elsewhere.




Figure 27: Strategy proportions averaged over different numbers of trials, varying the cost of selfishness (c_{self}). The 10,000th iteration averaged over a different numbers of trials, default parameters used elsewhere.

E. Varying b_{alt}




Figure 28: The average strategy proportions at each iteration, varying the benefit of altruism (b_{alt}). The average of 100 trials at each iteration, default parameters used elsewhere.




Figure 29: Strategy proportions averaged over different numbers of trials, varying the benefit of altruism (b_{alt}). The 10,000th iteration averaged over a different numbers of trials, default parameters used elsewhere.

F. Varying t




Figure 30: The average strategy proportions at each iteration, varying the time in pain (t). The average of 100 trials at each iteration, default parameters used elsewhere.




Figure 31: Strategy proportions averaged over different numbers of trials, varying the time in pain (t). The 10,000th iteration averaged over a different numbers of trials, default parameters used elsewhere.

G. The number of agents injured

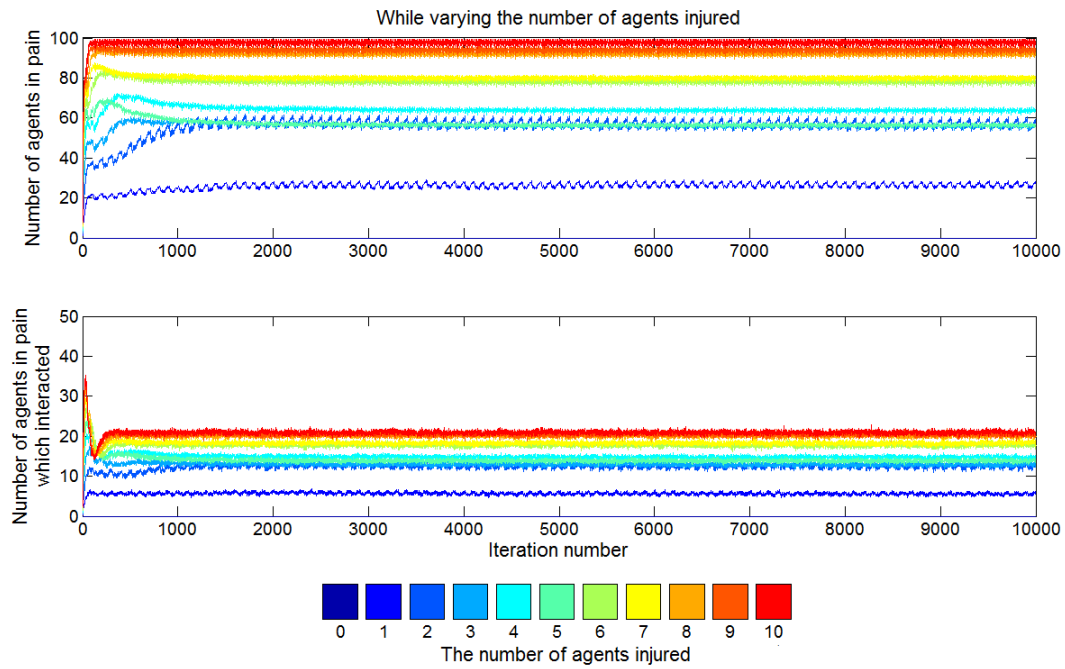


Figure 32: The effect of changing the number of agents injured each iteration (n) on the number of agents injured (above) and the number in interactions (below). Found using the average of 100 trials and default parameters.




Figure 33: The average strategy proportions at each iteration, varying the number of agents injured each iteration (n). The average of 100 trials used, default parameters used elsewhere.




Figure 34: Strategy proportions averaged over different numbers of trials, varying the interaction scaling. The 10,000th iteration averaged over a different numbers of trials, default parameters used elsewhere.

H. Sociability shift

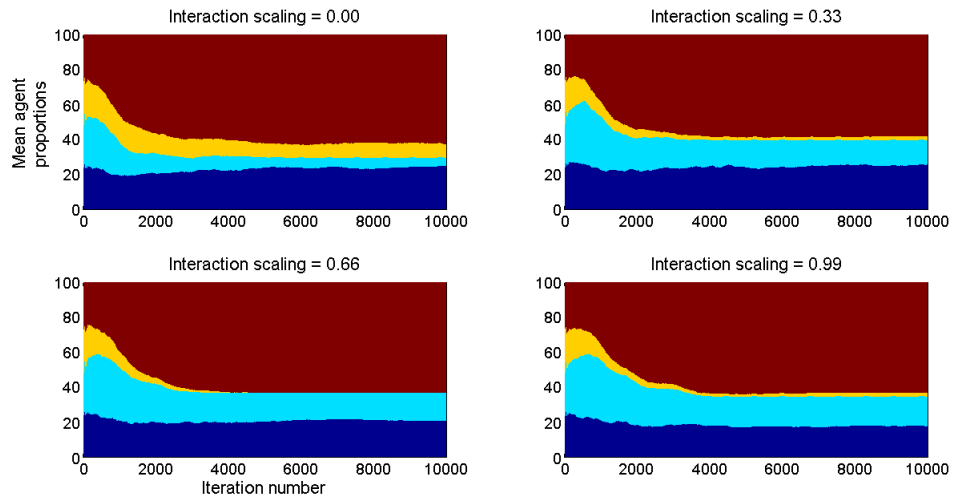


Figure 35: The average strategy proportions at each iteration, varying the shift in sociability. The average of 100 trials at each iteration, default parameters used elsewhere.

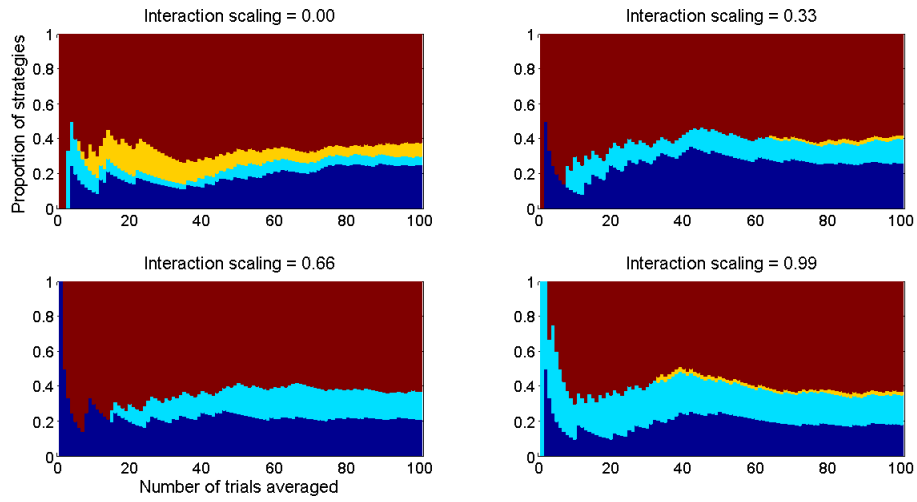


Figure 36: Strategy proportions averaged over different numbers of trials, varying the shift in sociability. The 10,000th iteration averaged over a different numbers of trials, default parameters used elsewhere.

I. Antagonists

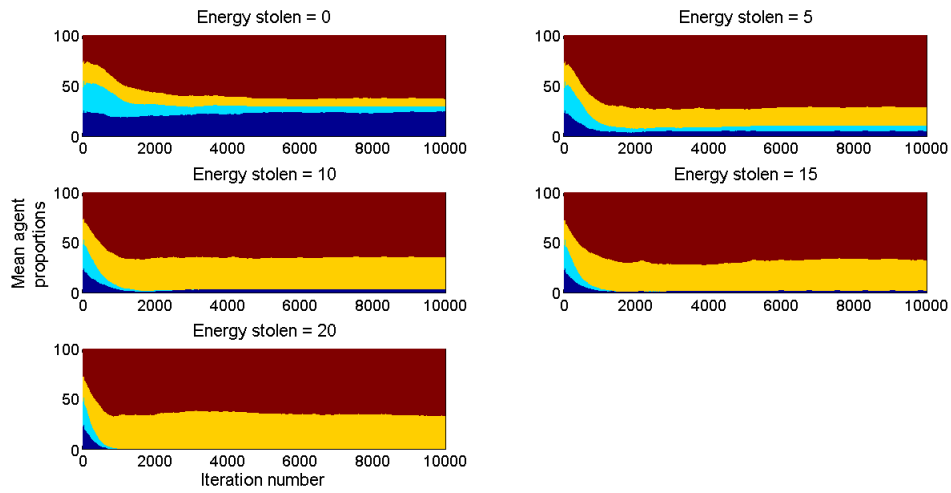


Figure 37: The average strategy proportions at each iteration, varying the amount of energy stolen in an environment with antagonists. The average of 100 trials at each iteration, default parameters used elsewhere.

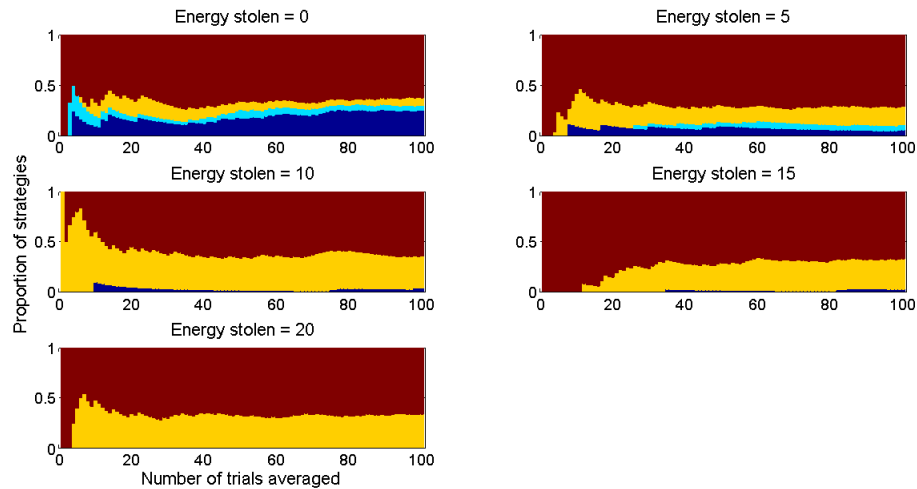


Figure 38: Strategy proportions averaged over different numbers of trials, varying the energy stolen in an environment with antagonists. The 10,000th iteration averaged over a different numbers of trials, default parameters used elsewhere.

J. Inheriting connectedness

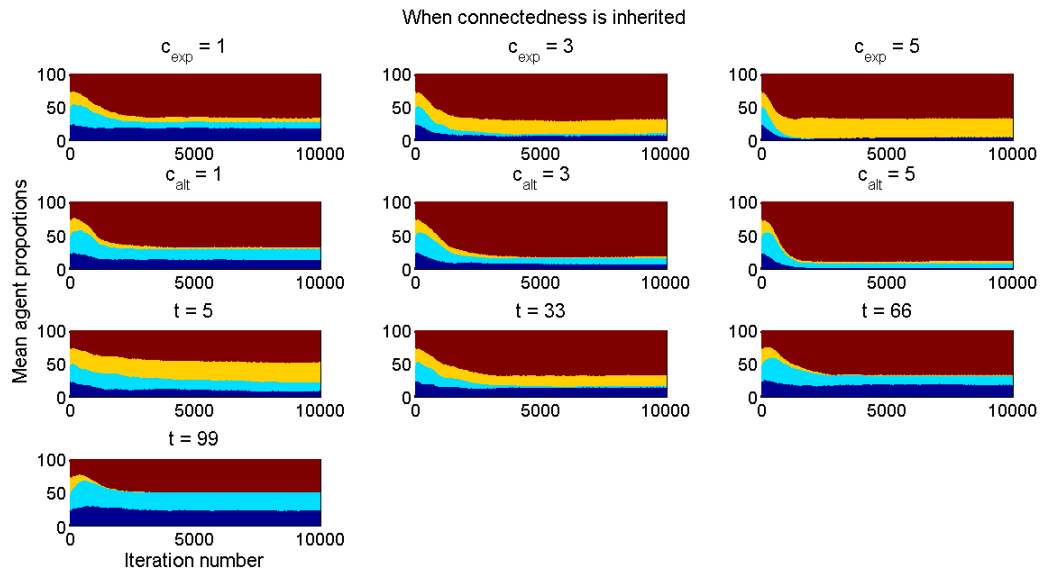


Figure 39: The average strategy proportions at each iteration, varying certain parameters in a model where connectedness can be inherited. The average of 100 trials at each iteration, default parameters used elsewhere.

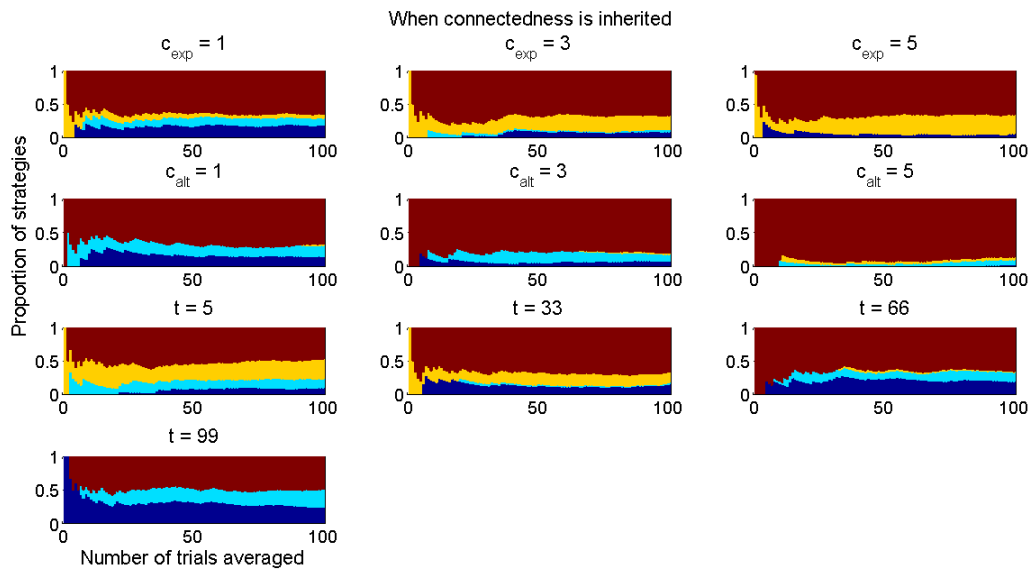


Figure 40: Strategy proportions averaged over different numbers of trials, varying certain parameters in a model where connectedness can be inherited. The 10,000th iteration averaged over a different numbers of trials, default parameters used elsewhere.

K. MATLAB code

Finding the initial agent population

```

1 function agents = ...
    initialagents(numberagents, minc, maxc, maxage, perexandalt, perexandself, ...
2         pernonexandalt, pernonexandself, initenergy)
3
4 % INPUT:
5
6 % numberagents:    number of agents wanted
7 % minc:           minimum connectedness score
8 % maxc:           maximum connectedness score
9 % maxage:         maximum age possible
10 % perexandalt:    initial proportions of altruistic expressors
11 % perexandself:   initial proportions of selfish expressors
12 % pernonexandalt: initial proportions of altruistic non-expressors
13 % pernonexandself: initial proportions of selfish non-expressors
14 % initenergy:     initial energy of agents
15
16 %% ===== CONNECTEDNESS =====
17
18 % generate 100 random elements from an exponential distribution with mean 1
19 meancon = 1;
20 alldist = exprnd(meancon, numberagents, 1);
21 % normalise to be between maxc and minc:
22 maxd = max(alldist);
23 normalised = (alldist/maxd) * (maxc - minc);
24 % connectedness values for intial agent population:
25 con = normalised + minc;
26
27 %% ===== AGE =====
28
29 % each agent can be an age between 0 and 100, we will start with a random
30 % distribution of these ages in our initial population
31 ages = round(rand(numberagents, 1) * (maxage - 1)); % so it can't be maxage ...
    when it starts
32
33 %% ===== INTIAL STRATEGIES =====
34
35 % first column: 1 if an expressor, 0 if a non-expressor
36 % second column: 1 if an altruist, 0 if selfish
37 e1 = [ones(perexandalt * numberagents, 2)];
38 e2 = [ones(perexandself * numberagents, 1), zeros(perexandself * numberagents, 1)];
39 e3 = ...
    [zeros(pernonexandalt * numberagents, 1), ones(pernonexandalt * numberagents, 1)];
40 e4 = [zeros(pernonexandself * numberagents, 2)];
41

```

```
42 ae = [e1;e2;e3;e4];
43 rae = ae(randperm(length(ae)),:);
44
45 %% ===== INITIAL AGENT POPULATION =====
46 agents = [rae(:,1) rae(:,2) con initenergy*ones(numberagents,1) ages ...
            zeros(numberagents,1) (1:numberagents)'];
47
48 end
```

Implementing the agent-based model

```

1 function [totstrats distconn disten meanenergystrats numdied numberinpain ...
   numberofinteractions agesoftheoneswhowilldie] = ...
2   painabm(totaliterations,maxc, minc,numberagents,initenergy,...
3   maxenergy,maxage,ageinc,forageinc,timepain,numinjured,inheritcon,...
4   cexp,calt,cself,balt,agents,interactionscaling)
5
6 % INPUT:
7 % totaliterations:  number of iterations
8 % maxc:             maximum connectedness
9 % minc:             minimum connectedness
10 % numberagents:    number of agents
11 % initenergy:      initial energy that all agents have
12 % maxenergy:       maximum energy that an agent can have
13 % maxage:          maximum age before an agent dies
14 % ageinc:          how much an agent ages each iteration
15 % forageinc:       how much energy received from foraging each iteration
16 % timepain:        maximum number of iterations in pain (if not helped)
17 % numinjured:     number of agents put into pain each iteration, n
18 % inheritcon:     1 or 0, whether you should inherit connectedness,
19 %                 or 0: just have a random connectedness
20 % cexp:            energy cost of expressing
21 % calt:            energy cost of helping an agent in pain
22 % cself:           connectedness cost of being selfish
23 % balt:            connectedness benefit of being altruistic
24 % agents:          the initial population of agents
25 % interactionscaling: sociability shift, s
26 %                 how much an agent's connectedness score is shifted when
27 %                 finding it's chance of interacting.
28 %
29 %
30 % OUTPUT:
31 % totstrats:       the number of agents using different strategies at each
32 %                 iteration
33 % distconn:        distribution of connectedness each iteration
34 % disten:          distribution of energy levels each iteration
35 % meanenergystrats: mean energy of agents using different strategies each
36 %                 iteration
37 % numdied:         total agents who got replaced each iteration
38 % numberofinteractions:
39 %                 number of interactions there were each iteration
40 % numberinpain:    total number in pain each iteration
41 % agesoftheoneswhowilldie: ages of the ones who have died each iteration.
42 %
43 bexp = timepain;   % benefit of expressing pain is that your time in pain ...
   goes to zero.
44 meancon = 1;      % the mean of the exponential distribution used to get ...
   connectedness values.
45
46 %% AN AGENT:
47 % [expressor?(1/0), altruistic?(1/0), connectedness score, energy, age, ...
   time-out, index number];

```

```

48
49 for t = 1:totaliterations
50     ttt = 1; % ages of the ones which will die counter
51
52     %% ===== CHOOSING AGENTS FOR INJURY =====
53
54     % which agents will be selected for pain this iteration?
55     % 'numinjured' random agents get injured (who are not already in pain)
56
57     if sum(agents(:,6)==0)≠0 % if there are healthy agents
58         if numinjured>size(agents(:,6)==0,1)
59             % if there are more chosen for injury than agents who are
60             % healthy, just injur all the remaining healthy agents
61             inpain = agents(:,7);
62         else
63             % pick 'numinjured' lots of random agents for injury
64             inpain = datasample(agents(agents(:,6)==0,7),numinjured);
65         end
66         %for these/this chosen agent the time in pain is set from 0 to ...
67         % "timepain"
68         agents(inpain,6) = timepain;
69     end
70     clear inpain
71
72     %% ===== CHOOSING INJURED AGENTS FOR INTERACTION =====
73
74     % Being interacted with is based on connectedness scores
75
76     %the agents in pain:
77     theagentsinpain = agents(agents(:,6)≠0,:);
78
79     if size(theagentsinpain,1)==1
80         % for the agents in pain, if there are any, find out if they get ...
81         % an interaction or
82         % not based on their connectedness score
83         ischosenint = zeros(size(theagentsinpain,1),1);
84         for i = 1:size(theagentsinpain,1)
85             % include sociability shift
86             bb = theagentsinpain(i,3)+interactionscaling;
87             if bb>maxc
88                 % if an agent has this shifted score > maxc, set it to maxc
89                 bb = maxc;
90             end
91             % find whether it interacts or not
92             if rand(1,1)<bb
93                 ischosenint(i) = theagentsinpain(i,7);
94             end
95             clear bb
96         end
97         % the indexes of the agents chosen for interaction:
98         chosenint = ischosenint(ischosenint≠0);
99     end
100     %% ===== BEGIN INTERACTIONS =====

```

```

101
102 % if there are interactions this iteration
103 if exist('chosenint')==1
104     ni = length(chosenint); %the number of agents being interacted with
105     % for these chosen agents in pain, let them interact with any ...
106     agent randomly
107     hchosenint = zeros(ni,1); % preallocation for pairing agents
108     for eachchosen = 1:ni
109         hchosenint(eachchosen) = ...
110             randsample(agents(agents(:,7)~=chosenint(eachchosen)...
111                 &~ismember(agents(:,7),hchosenint),7),1);
112     end
113
114 % interaction pairs
115 if size(chosenint,2)~=1
116     intpairs = [chosenint' hchosenint];
117 else
118     intpairs = [chosenint hchosenint];
119 end
120
121 % COSTS AND BENEFITS OF THE INTERACTIONS:
122 for i = 1:ni
123     if (agents(intpairs(i,1),1)==1)&&(agents(intpairs(i,2),2)==1) ...
124         % expressor vs altruistic
125
126         % AGENT IN PAIN, EXPRESSING:
127         agents(intpairs(i,1),4) = agents(intpairs(i,1),4) - cexp; ...
128         % energy cost of expressing
129         agents(intpairs(i,1),6) = agents(intpairs(i,1),6) - bexp; ...
130         % less iterations left in pain
131         % HEALTHY AGENT WHO IS INTERACTING, ALTRUISTIC:
132         agents(intpairs(i,2),4) = agents(intpairs(i,2),4) - calt; ...
133         % energy cost of helping
134         agents(intpairs(i,2),3) = agents(intpairs(i,2),3) + balt; ...
135         % connectedness increases
136
137     elseif ...
138         (agents(intpairs(i,1),1)==1)&&(agents(intpairs(i,2),2)==0) ...
139         % expressor vs selfish
140
141         % AGENT IN PAIN, EXPRESSING:
142         agents(intpairs(i,1),4) = agents(intpairs(i,1),4) - cexp; ...
143         % energy cost of expressing
144         % subtract the energy stolen here if using antagonists
145         % HEALTHY AGENT WHO IS INTERACTING, SELFISH:
146         agents(intpairs(i,2),3) = agents(intpairs(i,2),3) - cself; ...
147         % connectedness decreases
148         % add the energy stolen here if using antagonists
149
150     else
151         % in the cases of non-expressor vs altruistic and ...
152         non-expressor vs selfish
153         % there are no costs or benefits
154     end
155 end
156 end

```

```

144     end
145
146     %% ===== AGING, FORAGING AND PAIN RECOVERY=====
147
148     % let the agents age
149     agents(:,5) = agents(:,5) + ageinc;
150
151     % for the ones not in pain let them forage
152     notinpainagents = agents(agents(:,6)==0,7);
153     agents(notinpainagents,4) = agents(notinpainagents,4) + forageinc;
154
155     % if an agent in pain has not been helped in this iteration, decrease
156     % it's time left in pain by 1
157
158     for i = 1:numberagents
159         if agents(i,6)≠0
160             agents(i,6) = agents(i,6)-1;
161         end
162     end
163
164     %% ===== MAKE SURE LEVELS ARE WITHIN LIMITS =====
165
166     %connectedness:
167     agents(agents(:,3)>maxc,3) = maxc;
168     agents(agents(:,3)<minc,3) = minc;
169
170     %iterations left in pain:
171     agents(agents(:,6)>timepain,6) = timepain;
172     agents(agents(:,6)<0,6) = 0;
173
174     %energy:
175     agents(agents(:,4)>maxenergy,4) = maxenergy;
176     agents(agents(:,4)<0,4) = 0;
177
178     %age:
179     agents(agents(:,5)>maxage,5) = maxage;
180
181     %% ===== DEATH/BIRTH =====
182
183     % which agents should die now? death for: energy = 0 and age = maxage
184
185     j=1;
186     for i = 1:numberagents
187         if agents(i,4)==0||agents(i,5)==maxage; % the agents who should die.
188             todelete(j) = i; % indexes of those who will be deleted.
189             agesoftheoneswhowilldie(t,ttt) = agents(i,5);% what is their age?
190             ttt = ttt+1;
191             j = j+1;
192         end
193     end
194
195     if exist('todelete') % if there are agents due to die
196         %% ===== SELECTING PARENTS =====
197
198         % agents not about to be deleted:

```

```

199     healthyagents = agents(~ismember(agents(:,7),todelete),:); %all ...
        the details of these agents
200
201     % parents chosen weighted on their energy score
202     xp = healthyagents;
203     m = randn(1,1000);
204     mm = (1-abs(m)/max(abs(m)))*maxenergy;
205     % a normal distribution with values between 0 and 100
206
207     for i = 1:length(todelete)*2 % 2 parents per new agent
208         % normalise the possible parents' energy scores:
209         normhealthyagents = [xp(:,4)./(max(xp(:,4))) xp(:,7)];
210
211         % which of these possible parents' energy scores is closest to
212         % a number selected from mm
213         [cp indexp] = min(abs(normhealthyagents(:,1)-mm(i)));
214
215         % this agent will be chosen for parenthood
216         % if there is more than one with this same closest energy
217         % score, then pick randomly between them
218
219         if sum(abs(normhealthyagents(:,1)-mm(i))==cp)≠1
220             % if there is more than one agent with this same energy
221
222             rp = sum(abs(normhealthyagents(:,1)-mm(i))==cp); % number ...
                that are the same
223             rrp = randsample(1:rp,1); % pick one randomly
224             rrrp = ...
                normhealthyagents(abs(normhealthyagents(:,1)-mm(i))==cp,2);
225
226             % select this agent to be a parent
227             chosenparent(i) = rrrp(rrp);
228
229             % take this agent out of the list, to select another:
230             xp = xp(xp(:,7)≠rrrp(rrp),:);
231
232         else
233             % if there is only one agent with this energy
234             % select this agent for parenthood:
235             chosenparent(i) = normhealthyagents(indexp,2); % this is ...
                added to the list of agents to interact
236
237             % take this agent out of the list, to select another:
238             xp = xp(xp(:,7)≠normhealthyagents(indexp,2),:);
239         end
240         clear normhealthyagents cp indexp
241     end
242
243     %% ===== AGENT REPLACEMENT =====
244
245     % replace 'todelete' dying agents with an offspring agent from the
246     % selected parents
247
248     if exist('chosenparent')==1 % if there are any dying agents
249

```

```

250     % parent pairs
251     parents = reshape(chosenparent,length(todelete),2);
252
253     % distribution to take connectedness value from if inheritcon =
254     % 0
255     babyalldist = exprnd(meancon,1000,1);
256     babymaxd = max(babyalldist);
257     babynormalised = (babyalldist/babymaxd)*(maxc-minc);
258     babyconnectedness = babynormalised+minc;
259
260     for i = 1:length(todelete)
261
262         % connectedness
263         if inheritcon == 1 % inherit connectedness
264             babycon = ...
                mean([agents(parents(i,1),3),agents(parents(i,2),3)]);
265         else
266             babycon = babyconnectedness(i); % or select from ...
                exponential dist
267         end
268
269         % expression
270         if agents(parents(i,1),1)~=agents(parents(i,2),1)
271             babyexp = round(rand(1));
272         else
273             babyexp = agents(parents(i,1),1);
274         end
275
276         % altruism
277         if agents(parents(i,1),2)~=agents(parents(i,2),2)
278             babyalt = round(rand(1));
279         else
280             babyalt = agents(parents(i,1),2);
281         end
282
283         % replace the old agent with the new one, same index as ...
                the old one.
284         agents(todelete(i),:) = [babyexp babyalt babycon ...
                initenergy 0 0 todelete(i)]; % new agent
285         clear babycon babyexp babyalt
286     end
287 end
288 clear babyalldist babymaxd babynormalised babyconnectedness
289 end
290
291 %% ===== VARIABLES TO OUTPUT =====
292
293 % the number of agents using different strategies there are:
294 totstrats(1,t) = size(agents((agents(:,1)==1)&(agents(:,2)==1),:),1);
295 totstrats(2,t) = size(agents((agents(:,1)==1)&(agents(:,2)==0),:),1);
296 totstrats(3,t) = size(agents((agents(:,1)==0)&(agents(:,2)==1),:),1);
297 totstrats(4,t) = size(agents((agents(:,1)==0)&(agents(:,2)==0),:),1);
298
299 if size(agents,1)==numberagents
300     distconn(t,1:numberagents) = agents(:,3);

```

```

301     disten(t,1:numberagents) = agents(:,4);
302 else
303     distconn(t,1:numberagents) = zeros(1,numberagents);
304     disten(t,1:numberagents) = zeros(1,numberagents);
305 end
306
307 % mean energy for each strategy
308 meanenergystrats(1,t) = mean(agents(agents(:,1)==1&agents(:,2)==1,4));
309 meanenergystrats(2,t) = mean(agents(agents(:,1)==1&agents(:,2)==0,4));
310 meanenergystrats(3,t) = mean(agents(agents(:,1)==0&agents(:,2)==1,4));
311 meanenergystrats(4,t) = mean(agents(agents(:,1)==0&agents(:,2)==0,4));
312
313 if exist('todelete')
314     numdied(t) = length(todelete); % total agents who got replaced
315 else
316     numdied(t) = 0;
317 end
318
319
320 if exist('ni')
321     numberofinteractions(t) = ni; % number of interactions there were ...
        this iteration
322 else
323     numberofinteractions(t) = 0;
324 end
325
326 numberinpain(t) = size(agents(agents(:,6)~=0),1);
327
328 clear ww numhealthy numinterations numpain yy yyy agentsinpain x ...
        chosenint ...
329     ni intpairs todelete healthyparents xp m mm chosenparent parents ...
        theagentsinpain ischosenint ttt
330 end
331 end

```


References

- Axelrod, R. M. 1997. *The complexity of cooperation: Agent-based models of competition and collaboration*. Princeton University Press.
- Axelrod, R., & Dion, D. 1988. The further evolution of cooperation. *Science* 242(4884): 1385-1390.
- Barrier, A. C., Ruelle, E., Haskell, M. J., & Dwyer, C. M. 2012. Effect of a difficult calving on the vigour of the calf, the onset of maternal behaviour, and some behavioural indicators of pain in the dam. *Preventive Veterinary Medicine* 103(4): 248-256.
- Black, H. 2009. Underwater Suffering? *Scientific American Mind* 20(5): 16-16.
- Block, A. R., Kremer, E. F., & Gaylor, M. 1980. Behavioral treatment of chronic pain: the spouse as a discriminative cue for pain behavior. *Pain* 9(2): 243-252.
- Bolles, R. C., & Fanselow, M. S. 1980. A perceptual-defensive-recuperative model of fear and pain. *Behavioral and Brain Sciences* 3(2): 291-301.
- Brown, J. L., Sheffield, D., Leary, M. R., & Robinson, M. E. 2003. Social support and experimental pain. *Psychosomatic Medicine* 65(2): 276-283.
- Butler, R. K., & Finn, D. P. 2009. Stress-induced analgesia. *Progress in neurobiology* 88(3): 184-202.
- Coan, J. A., Schaefer, H. S., & Davidson, R. J. 2006. Lending a Hand Social Regulation of the Neural Response to Threat. *Psychological Science* 17(12): 1032-1039.
- Damasio, A. R. 1999. *The feeling of what happens: Body and emotion in the making of consciousness*.
- Epstein, J. M., & Axtell, R. 1996. *Growing artificial societies: social science from*

- the bottom up. *Brookings Inst. Press: Washington, DC*.
- Farmer, J. D., & Foley, D. 2009. The economy needs agent-based modelling. *Nature* 460(7256): 685-686.
- Fordyce, W. E. 1976. *Behavioral methods for chronic pain and illness* pp. 86-90. St. Louis: Mosby.
- Gilbert, N. 2008. Agent-based models (No. 153). Sage.
- Grunau, R. V., & Craig, K. D. 1987. Pain expression in neonates: facial action and cry. *Pain* 28(3): 395-410.
- Hogeweg, P., & Hesper, B. 1983. The ontogeny of the interaction structure in bumble bee colonies: a MIRROR model. *Behavioral Ecology and Sociobiology* 12(4): 271-283.
- Keating, S. C., Thomas, A. A., Flecknell, P. A., & Leach, M. C. 2012. Evaluation of EMLA cream for preventing pain during tattooing of rabbits: Changes in physiological, behavioural and facial expression responses. *PloS one* 7(9): e44437.
- Langford, D. J., Cragger, S. E., Shehzad, Z., Smith, S. B., Sotocinal, S. G., Levenstadt, J. S., ... & Mogil, J. S. 2006. Social modulation of pain as evidence for empathy in mice. *Science* 312(5782): 1967-1970.
- Langford, D. J., Bailey, A. L., Chanda, M. L., Clarke, S. E., Drummond, T. E., Echols, S., ... & Mogil, J. S. 2010. Coding of facial expressions of pain in the laboratory mouse. *Nature methods* 7(6): 447-449.
- Langford, D. J., Tuttle, A. H., Briscoe, C., Harvey-Lewis, C., Baran, I., Gleeson, P., ... & Mogil, J. S. 2011. Varying perceived social threat modulates pain behavior in male mice. *The Journal of Pain* 12(1): 125-132.
- Macal, C. M., & North, M. J. 2006. Tutorial on agent-based modeling and simula-

- tion part 2: how to model with agents. In *Proceedings of the 38th Winter Simulation Conference, Monterey, CA*: 73-83.
- MacIntyre, L. C., Langford, D. J., Lacroix-Fralish, M. L., Austin, J-S., Sotocinal, S. G., Sternberg, W. F. & Mogil, J. S. 2007. Sex-dependent fear of human experimenters and facsimiles thereof in laboratory mice. . *Soc. Neurosci. Abstr.* 33.
- Mogil, J. S. 2009. Animal models of pain: progress and challenges. *Nature Reviews Neuroscience* 10(4): 283-294.
- Niazi, M., & Hussain, A. 2011. Agent-based computing from multi-agent systems to agent-based models: a visual survey. *Scientometrics* 89(2): 479-499.
- Nowak, M. A. 2006. Five rules for the evolution of cooperation. *Science* 314(5805): 1560-1563.
- Nowak, M. A., Bonhoeffer, S., & May, R. M. 1994. More spatial games. *International Journal of Bifurcation and Chaos* 4(01): 33-56.
- Peeters, P. A., & Vlaeyen, J. W. 2011. Feeling more pain, yet showing less: the influence of social threat on pain. *The Journal of Pain* 12(12): 1255-1261.
- Ranger, M., Johnston, C. C., & Anand, K. J. S. 2007, October. Current controversies regarding pain assessment in neonates. In *Seminars in perinatology* (Vol. 31, No. 5, pp. 283-288). WB Saunders.
- Reynolds, C. W. 1987, August. Flocks, herds and schools: A distributed behavioral model. In ACM SIGGRAPH *Computer Graphics* 21(4): 25-34. ACM.
- Schelling, T. C. 1971. Dynamic models of segregation. *Journal of mathematical sociology* 1(2): 143-186.
- Sullivan, M. J., Thibault, P., Savard, A., Catchlove, R., Kozey, J., & Stanish, W. D. 2006. The influence of communication goals and physical demands on different

- dimensions of pain behavior. *Pain* 125(3): 270-277.
- Trivers, R. L. 1971. The evolution of reciprocal altruism. *Quarterly review of biology* 35-57.
- Wall, P. D. 1979. On the relation of injury to pain the John J. Bonica Lecture. *Pain* 6(3): 253-264.
- Wiech, K., & Tracey, I. 2013. Pain, decisions, and actions: a motivational perspective. *Frontiers in neuroscience* 7.
- Williams, A. C. D. C. 2002. Facial expression of pain: an evolutionary account. *Behavioral and brain sciences* 25(4): 439-455.
- Williams, A. C. D. C., & Craig, K. D. 2006. A science of pain expression?. *Pain* 125(3): 202-203.