

Fostering Co-operative Behaviour Through Social Intervention

Martyn Lloyd-Kelly¹, Katie Atkinson² and Trevor Bench-Capon²

¹*Institute of Psychology, Health and Society, University of Liverpool, Liverpool, U.K.*

²*Department of Computer Science, University of Liverpool, Liverpool, U.K.*

Keywords: Co-operation, Emotion, Agents, Simulation.

Abstract: The emergence and maintenance of co-operation in a society of agents is an important issue and some recent research has explored the role that can be played by a functional model of emotions. For example, it has been shown that the emotions of *gratitude* and *anger* can be used to produce co-operative behaviour in a public goods game from agents acting solely in accordance with their current emotional state. The effectiveness of these emotions in producing co-operation depends on the emotional volatility of the agents, which determines the strength of these emotions required to alter behaviour. Often, however, dysfunctional relationships develop, which impairs the performance of the society as a whole. In this paper we explore through a series of computational simulations how interventions by society can be used to correct dysfunctional behaviour. The results of our simulations show that enforcement of co-operative behaviour and education to alter emotional characters can improve overall performance in the dysfunctional cases and that different interventions are appropriate given different initial circumstances.

1 INTRODUCTION

Emotions such as *gratitude* and *anger* have been shown to be key determinants of social behaviour in humans: gratitude is an essential motivator of co-operative and pro-social behaviour (Vohs et al., 2006; DeSteno et al., 2010), whilst altruistic punishment resulting from anger is another essential element in sustaining human co-operation (Fehr and Gächter, 2002; Hirshleifer, 1987). Several papers provide logical formalisations of emotions for agent systems and demonstrate how emotions can be computationally modelled to play a functional and beneficial role in enabling agents to determine how to respond to environmental information (Steunebrink et al., 2007; Adam et al., 2009).

Using computational simulation, emotions have been used to determine responses for agents playing iterated *Prisoner's Dilemma* games (Bazzan and Bordini, 2001; Lloyd-Kelly et al., 2012a; Lloyd-Kelly et al., 2012b). The latter two papers show how the emergent property of co-operation, achievable by game theoretic behaviour in the context of the *Prisoner's Dilemma* (Axelrod, 1984), can also be achieved using simulated emotions. Our aim in this work is to explore the *effect* of emotions and different *emotional characters* on fostering co-operation

through computational simulation using the iterated Prisoner's Dilemma as a test-bed.

Interactions between emotional agents playing iterated Prisoner's Dilemma games tend to go through a period of initial variation and then settle down to one of four relationships: *mutual co-operation*, *mutual defection*, *symmetric turn-taking* or *asymmetric turn-taking* (Lloyd-Kelly et al., 2012a; Lloyd-Kelly et al., 2012b). In symmetric turn-taking, an agent, x , co-operates in round n while its opponent, y , defects in n . This continues for m rounds until, in round $n + m$, the behaviour of x and y switches (in this example, y co-operates and x defects). This continues for l rounds before behaviour switches back on round $n + m + l$. If $m = l$ turn-taking is symmetric; agents co-operate and defect for an equal number of rounds before switching, otherwise it is asymmetric; one agent defects for more rounds before switching behaviour. Over time, symmetric turn-taking tends to equality for the agents involved whilst both symmetric and asymmetric turn-taking are significantly better than mutual defection, from the point of view of the system. Asymmetric turn-taking is always better than mutual defection for the individual that defects more and is often better for the other agent also¹.

¹See table 1: defecting even once every four rounds secures greater individual utility rather than mutual defection.

Table 1: Prisoner’s Dilemma pay-off matrix.

		Player i	
		Co-op	Defect
Player j	Co-op	$3_i, 3_j$	$5_i, 0_j$
	Defect	$0_i, 5_j$	$1_i, 1_j$

In human society, the temptation to secure an individual advantage through defection when mutual co-operation exists is ubiquitous (Simpson, 2003) so, in this paper, *hope* is introduced as an emotion. Whether this emotion will lead to defection will depend on how likely it is that hope’s effect is manifest (here represented as a probability of defection when hope is active). Consequently, hope potentially destabilises mutual co-operation and the defecting agent may gain some advantage (which is what is hoped for) before the relationship stabilises again.

In the context of a society playing the Prisoner’s Dilemma, mutual co-operation is best since it promotes greater aggregate scores. Turn-taking (symmetric or asymmetric), may also be acceptable since the reduction of aggregate score is relatively small. Mutual defection is, however, dysfunctional since it scores only a third of that achieved by mutual co-operation, and 40% of that achieved by turn-taking. In this paper we will explore the effects of two mechanisms implemented by society to intervene and attempt to correct dysfunctional behaviour: *enforcing co-operative behaviour* and altering the emotional characters of the agents concerned through *education*. Real examples of the latter are enforced attendance at speed awareness courses for minor speeding offences, or more generally the inclusion of rehabilitation in prison regimes.

The paper is structured as follows: section 2 describes how we model emotions whilst section 3 discusses the simulation environment used. Section 4 gives key results from simulations where emotional agents are not subject to intervention or education. Section 5 presents results of enforcing co-operative behaviour and applying to the simulations the four education strategies modelled. Section 6 offers a discussion of some interesting results, and section 7 presents the main conclusions.

2 MODELLING EMOTION

The emotion modelling framework used in this paper has been used in previous work (Lloyd-Kelly et al., 2012a; Lloyd-Kelly et al., 2012b) and is underpinned by the OCC model of emotion (Ortony et al., 1988). Use of the OCC model by computer scientists for

computationally modelling emotion (Bazzan and Bordini, 2001; Burghouts et al., 2003; Steunebrink et al., 2007; Adam et al., 2009) is largely due to its tractability and use of concepts that are closely linked to those used by the agent systems community. The OCC defines emotions as valenced reactions by an agent to the consequences of an event (relevant to the agent itself or another), an action (performed by the agent itself or another) or some aspect of an object in relation to the agent’s current goals. Numeric emotion *potentials* are increased or decreased as the agent’s goals are impacted by the three factors mentioned. When an agent’s emotion potential equals the emotion’s activation threshold, the agent’s intentional behaviour is affected. Therefore, an agent’s intentional behaviour is determined by its current emotional state which itself is a product of all interactions the agent has had with its opponent in accordance with the descriptions of these emotions given in the OCC model.

Nine emotional characters are defined in our previous work (Lloyd-Kelly et al., 2012a; Lloyd-Kelly et al., 2012b); different emotional characters have different activation thresholds for emotions (see table 2). The emotional character notation used in table 2 is based upon the first letter of an emotion and the activation thresholds modelled for that emotion². The difference in activation thresholds for anger and gratitude result in degrees of *tolerance* and *responsiveness*: in table 2 tolerance increases down the rows and responsiveness decreases across the columns (from left to right). Tolerance is determined by anger’s activation threshold and indicates the readiness of an agent to punish (the higher the activation threshold for anger, the slower an emotional character is to punish and the more tolerant it is said to be). Responsiveness is determined by gratitude’s activation threshold and indicates the readiness of an agent to reward (the higher the activation threshold for gratitude, the slower an emotional character is to reward and the less responsive it is said to be). Computational simulations show that high tolerance and responsiveness promote total system scores when agents play iterated Prisoner’s Dilemma games (Lloyd-Kelly et al., 2012b). This is because a high tolerance and responsiveness increases the likelihood of establishing and maintaining co-operation, even when the opponent initially defects.

In this paper, the nine emotional characters described in table 2 are modelled and augmented with hope, resulting in three emotions being modelled: anger, gratitude and hope. These emotions will be used to explore the interaction between emotional

²For example, the activation thresholds for Anger and Gratitude in A1:G1 are both set to 1.

Table 2: Emotional character descriptions.

		If defecting, #co-ops required to co-op.		
		1	2	3
If co-op, #defects required to defect.	1	A1:G1	A1:G2	A1:G3
	2	A2:G1	A2:G2	A2:G3
	3	A3:G1	A3:G2	A3:G3

Table 3: Details of emotions modelled.

Emo.	Eli. Cond.	Act. Thresh.	Effect	Prob. Effect.
Anger	Opp. defect	1/2/3	Defect	1
Gratitude	Opp. co-op	1/2/3	Co-op	1
Hope	Mutual co-op	3	Defect	0.4

character and societal intervention. The emotion modelling framework used here stipulates that an emotion must have the following factors defined: eliciting conditions, potential, activation threshold, saturation, effect and probability of effect (Lloyd-Kelly, 2014). These factors (excluding emotion potential since this is variable but is initialised to 0 and saturation which is always equal to the activation threshold) for each emotion are given in table 3. When an eliciting condition for an emotion is registered by an agent, the emotion’s potential increases by 1 until its saturation point is reached. When an emotion’s potential is greater than or equal to its activation threshold, the emotion’s effect has a probability (0-1) of being manifest in the agent.

Justifications for anger and gratitude’s eliciting conditions, effects and probability of effect are provided in previous work (Lloyd-Kelly et al., 2012a; Lloyd-Kelly et al., 2012b); eliciting conditions and effects of hope are based upon work by Lazarus and Snyder (Lazarus, 1999; Snyder, 2002). Lazarus states that hope is elicited by a strong desire to be in a situation that is more favourable than the current one and Snyder states that hope motivates agents to derive pathways to desired goals and use agency thinking to traverse those pathways. Therefore, hope causes the agent to defect with some probability. We use 0.4, a mid-range value for which co-operation is more likely. A defecting agent will obtain 5 points rather than 3 if the opponent continues to co-operate. Algorithms for anger, gratitude and hope are provided in sections 2.1 and 2.2 for clarification. Note that emotions are directed towards specific agents so an agent, x , may be angry with its opponent, y , and not hopeful

of defecting against it whilst simultaneously grateful to its opponent, z , and hopeful of defecting against it.

2.1 Anger and Gratitude Algorithm

1. Agent x analyses its opponent y ’s behaviour in the round just played.
 - (a) If y defected and x ’s anger potential for y is $<$ x ’s anger saturation value, x increases its anger potential for y by 1.
 - (b) If y co-operated and x ’s gratitude potential for y is $<$ x ’s gratitude saturation value, x increases its gratitude potential for y by 1.
2. Agent x checks anger and gratitude potentials for y .
 - (a) If x ’s anger potential for $y \geq x$ ’s anger activation threshold, x ’s behaviour towards y is set to defect and x ’s gratitude potential for y is reset to 0.
 - (b) If x ’s gratitude potential for $y \geq x$ ’s gratitude activation threshold for y , x ’s behaviour towards y is set to co-operate and x ’s anger potential for y is reset to 0.
 - (c) If x ’s anger or gratitude potentials for y are $<$ their activation thresholds, x ’s behaviour towards y is determined by x ’s initial behaviour setting.

2.2 Hope Algorithm

1. Agent x analyses its opponent y ’s behaviour in the round just played.
 - (a) If y defected, x ’s hope potential for y is reset to 0.
 - (b) If y co-operated, x checks own behaviour in the round just played:
 - i. If x also co-operated and x ’s hope potential for y is $<$ x ’s hope saturation value, increase hope potential for y by 1.
 - ii. If x defected, reset hope potential for y to 0.
2. Agent x checks hope potential towards y . If x ’s hope potential for $y \geq x$ ’s hope activation threshold, x randomly selects an integer, i between 1 and 10:
 - (a) If $i \leq 4$, x will defect against y in the next round.
 - (b) If $i > 4$, x will defect/co-operate with y in the next round according to x ’s anger/gratitude/initial behaviour status.

3 SIMULATION DETAILS

The simulation environment used is based upon that used in previous work (Lloyd-Kelly et al., 2012a; Lloyd-Kelly et al., 2012b): 576 agents play in context of an iterated Prisoner’s Dilemma game on a virtual, two-dimensional, 24×24 , grid whose edges are not bound. Agents remain stationary and always play against their 8 immediate neighbours at their cardinal and inter-cardinal compass points. Since there are nine emotional characters, 576 agents ensures that emotional characters are equally represented (64 agents of each emotional character) in mixed population simulations. The iterated Prisoner’s Dilemma game permits agents to evaluate events and modify their emotional state and intentional behaviour in light of them³. Furthermore, there are very few variables that can influence outcomes, ensuring that any results are relatively free from contamination by unmanaged variables. Finally, since agents are only capable of either co-operating or defecting, this facilitates a determination of emotions to be modelled along with their eliciting conditions and effects.

The pay-off distribution for the standard Prisoner’s Dilemma formalisation (see table 1) is used. The number of agents whose initial behaviour is set to co-operate/defect for each scenario run is kept consistently equal (288) and the location of agents within the simulation environment is randomised to ensure no experimenter bias with respect to spatial distribution of emotional characters or initial behaviour which could skew results.

In section 1, we mentioned that agents will be subject to intervention from society through behaviour enforcement or education; table 4 indicates which simulations use behaviour enforcement and education strategies. Behaviour enforcement occurs when an agent and its opponent have mutually defected 3 times consecutively and immediately results in the re-establishment of mutual co-operation and all emotion potentials being reset to 0 for these agents. We chose 3 mutual defections as a condition for behaviour enforcement since none of the emotional characters implemented are capable of re-establishing co-operation with an opponent after 3 mutual-defections by virtue of their emotional character. If education is used, the agent’s anger or gratitude threshold is additionally incremented or decremented by 1⁴. With our emotional characters there are four possible education strategies:

³This is obviously not possible using the “one-shot” variation of the game.

⁴Agents will only have their anger/gratitude thresholds decremented until they equal 1; the minimum value for these thresholds.

Table 4: Emotions, behavioural enforcement and education strategy used in each simulation type.

Sim. Type	Emotions	Behv. Enf.	Edu. Str.
1	A, G	-	-
2	A, G, H	-	-
3	A, G, H	✓	-
4	A, G, H	✓	1
5	A, G, H	✓	2
6	A, G, H	✓	3
7	A, G, H	✓	4

1. **Increase tolerance.** Add 1 to an agent’s anger activation threshold. This can be seen as *anger management* since the agent will not punish others so quickly.
2. **Decrease responsiveness.** Add 1 to the agent’s gratitude activation threshold. This represents a form of *assertiveness training* where the agent is less likely to be misled by co-operative behaviour and is made more ready to insist on its rights.
3. **Decrease tolerance.** Subtract 1 from an agent’s anger activation threshold. This can be seen as another form of *assertiveness training*, in which defections are more likely to be punished.
4. **Increase responsiveness.** Subtract 1 from the agent’s gratitude activation threshold. This is analogous to increasing *empathy* and *altruism*, so that an agent is less suspicious of others and readier to respond to co-operation.

In total, 140 simulations were run; 70 for 500 rounds and 70 for 1000 rounds. Longer round-limits were used to ascertain the effects of education and intervention over longer periods of time. In this paper we will discuss the 500 round simulations in detail and comment on the effects of additional rounds in section 6. Each half of the 140 simulations consists of 7 simulation types which prescribe the emotions and intervention used. A succinct description of these simulation types can be found in table 4. For each simulation type, 10 emotional character populations were used: 1-9 are initially composed of a single emotional character whilst population 10 is composed of an equal mix of all emotional characters. Note that initial behaviour setting for population 10 is not decided at the level of emotional characters but rather at the level of the population.

When the round limit is reached, the simulation environment is completely reset (including all variables for all agents) and the total system score (the sum of every agent’s final score in the population) is recorded with the number of interventions that have occurred (if any).

Table 5: Total system scores for emotional character populations when hope is/is not modelled and loss incurred when hope is introduced.

Ch.	No Hope	Hope	Loss
A1:G1	5193000	4426616	766384
A1:G2	3463444	2316024	1147420
A1:G3	3437522	2316018	1121504
A2:G1	5744862	4709929	1034933
A2:G2	5196000	4517117	678883
A2:G3	3462912	2333223	1129689
A3:G1	5760846	4613945	1146901
A3:G2	5771668	4634982	1136686
A3:G3	5186000	4540339	645661
Mixed	4630504	2928854	1701650

4 WITHOUT INTERVENTION

Table 5 shows the total score for all 10 populations without behaviour enforcement or education. The emotional characters can be divided into two broad groups: one (A1:G2, A1:G3 and A2:G3) contains those that are quicker to punish than to respond to co-operation (and which perform relatively badly) and the other is composed of the remaining emotional characters from table 2 that are at least as ready to respond as to punish (and which perform relatively well). A mixed population performs somewhere between these two groups. Table 5 also shows how the introduction of opportunistic defection due to hope significantly decreases total system scores. Emotional characters that performed worse initially suffer more than the better performing emotional characters. Those which are equally ready to punish and to respond (A1:G1, A2:G2 and A3:G3) suffer least. The mixed group suffers worst of all.

5 WITH INTERVENTION

As mentioned in section 1, the simulations will settle down into mutual defection, turn-taking (whether symmetric or asymmetric), or periods of mutual co-operation (which break down when an agent succumbs to the temptation to defect due to hope). From the viewpoint of society as whole, mutual defection is highly detrimental to the system score and should be avoided. Therefore in this section we will explore the effects of intervention, described in section 3, designed to prevent mutual defection.

5.1 Enforcing Behaviour

Table 6 shows the effect of enforcing mutual co-operation following 3 rounds of mutual defection on

Table 6: Total system scores for hopeful emotional character populations when behaviour enforcement occurs without education.

Ch.	Score	Gain	# Int.	Gain per. Int.
A1:G1	5746674	1320058	2560	515.65
A1:G2	4297626	1981602	292636	6.77
A1:G3	4298293	1982275	292574	6.78
A2:G1	5487839	777910	1134	685.99
A2:G2	5739134	1222017	3126	390.92
A2:G3	4919034	2585811	138944	18.61
A3:G1	5387322	773377	1132	683.20
A3:G2	5415490	780508	1156	675.18
A3:G3	5731018	1190679	2826	421.33
Mixed	5106992	2178138	99502	21.89

the populations of table 5, giving the total score, the gain in total score achieved by each population when its total score is compared to the total score achieved by the same population when intervention does not occur and hope is present in the population (see table 5), the number of interventions made, and the gain per intervention.

From the results shown in table 6, we can see several points of interest. In all cases the intervention does much to remedy hope's effect. Indeed, except for the emotional characters which are more responsive than tolerant (A2:G1, A3:G1 and A3:G2), intervention improves the system score compared with the situation where hope is not present. The number of interventions required, however, varies substantially, with the dysfunctional emotional characters (A1:G2, A1:G3 and A2:G3) requiring a very large number of interventions, as they frequently return to mutual defection. The mixed population also requires very frequent intervention. It is the emotional characters which are more responsive than tolerant that require fewest interventions. Thus the gain per intervention is very small for emotional characters A1:G2, A1:G3 and A2:G3, and rather small for the mixed population: in these cases the gains may well not justify the cost of intervention. Other emotional characters attain substantial gains per intervention, especially A2:G1, A3:G1 and A3:G2, for which the smaller number of interventions means that the value for money is greatest for these emotional characters.

The improved score should exceed the cost of interventions, and so we can establish a maximum intervention cost⁵ from the final column of table 6: if the A1:G1 population expects to gain 515.65 per intervention then interventions should not cost more than this. Although intervention to enforce co-operation

⁵The maximum cost that can be paid to intervene without the total system score incurring a loss as a result.

Table 7: Total system scores for hopeful emotional character populations with behaviour enforcement and education strategy 1 (anger management).

Ch.	Score	Gain	# Int.	Gain per. Int.	Edu. Gain
A1:G1	5291340	864724	1134	762.54	-455334
A1:G2	5281348	2965324	3416	868.07	983722
A1:G3	5281345	2965327	3416	868.07	983052
A2:G1	5290610	580681	1160	500.59	-197229
A2:G2	5290353	773236	1134	681.87	-448781
A2:G3	5292018	2958795	1162	2546.30	372984
A3:G1	5288354	674409	1106	609.77	-98968
A3:G2	5290109	655127	1150	569.68	-125381
A3:G3	5290230	749891	1176	637.66	-440788
Mixed	5288846	2359992	1686	1399.76	181854

will improve the aggregate system scores, mutual defection will still develop as the enforced co-operation breaks down in the face of temptation. We assume that the resources to supervise every move are not available: the intervention should encourage agents to work together in a productive relationship.

5.2 Education

As argued in section 5.1, enforcing co-operation will not prevent mutual defection from developing. We now explore the use of education to alter emotional character, in the hope that this will reduce the recurrence of mutual defection.

Tables 7-10 show the results of the four different educational strategies on the various populations, together with the number of interventions and the average gain per intervention. The gain from education, rather than simply enforcing co-operation of dysfunctional pairs is also shown.

Table 7 shows that education strategy 1 has a dramatic effect on the number of interventions, so that all initial emotional characters now have an acceptable number of interventions, although emotional characters A1:G2 and A1:G3 require more than the rest. The gain per intervention is now uniformly good, and spectacular in the case of emotional character A2:G3. The mixed population also benefits considerably, especially from the dramatically reduced number of interventions, so that the gain per intervention is second only to emotion character A2:G3.

This education strategy is, however, not beneficial in terms of system score for emotional characters A1:G1, A2:G1, A2:G2, A3:G1, A3:G2 and A3:G3. Moreover, while education does have value for emotional characters A1:G1, A2:G2 and A3:G3, for those that required fewest interventions without education,

Table 8: Total system scores for hopeful emotional character populations with behaviour enforcement and education strategy 2 (assertiveness training 1).

Ch.	Score	Gain	# Int.	Gain per. Int.	Edu. Gain
A1:G1	4298740	-127876	292326	-0.44	-1447934
A1:G2	4296902	1980878	292720	6.77	-724
A1:G3	4297804	1981786	292649	6.77	-489
A2:G1	4714455	4526	196271	0.02	-773384
A2:G2	4708261	191144	197526	0.97	-1030873
A2:G3	4709001	2375778	197385	12.04	-210033
A3:G1	4885127	271182	150845	1.80	-502195
A3:G2	4878522	243540	152240	1.60	-536968
A3:G3	4879912	339573	152166	2.23	-851106
Mixed	4646843	1717989	211521	8.12	-460149

Table 9: Total system scores for hopeful emotional character populations with behaviour enforcement and education strategy 3 (assertiveness training 2).

Ch.	Score	Gain	# Int.	Gain per. Int.	Edu. Gain
A1:G1	5747270	1320654	2443	540.59	596
A1:G2	4297206	1981182	292608	6.77	-420
A1:G3	4296484	1980466	292710	6.77	-1809
A2:G1	5738050	1028121	4022	255.62	250211
A2:G2	4298235	-218882	292267	-0.75	-1440899
A2:G3	4300303	1967080	292297	6.73	-618731
A3:G1	5739073	1125128	3843	292.77	351751
A3:G2	4301740	-333242	291950	-1.14	-1113750
A3:G3	4296964	-243375	292458	-0.83	-1434054
Mixed	4472148	1543294	257417	6.00	-634844

the emotional characters which are initially more responsive than tolerant, the gain per intervention is reduced, so that for these emotional characters this form of education is not worthwhile.

From table 8 we see that the effects of education strategy 2, which is designed to reduce responsiveness, is in complete contrast to education strategy 1. In no case is education strategy 2 beneficial when compared with enforcing co-operation, and although there are gains from intervention for all emotional characters (except A1:G1), the number of interventions required as a result of applying education strategy 2 is considerably less value for money compared to when behaviour enforcement is the sole intervention used.

Table 9 gives a similar picture to table 8. Characters A1:G1, A2:G1 and A3:G1 do show some gains from education, but these are small, and are outweighed by the increased number of interventions required. Therefore education strategy 3 is not recommended, unless the desire to maximise overall sys-

Table 10: Total system scores for hopeful emotional character populations with behaviour enforcement and education strategy 4 (increase empathy/altruism).

Ch.	Score	Gain	# Int.	Gain per. Int.	Edu. Gain
A1:G1	5746418	1319802	2608	506.06	-256
A1:G2	5728120	3412096	6256	545.41	1430494
A1:G3	5728255	3412237	6243	546.57	1429962
A2:G1	5487031	777102	1144	679.28	-808
A2:G2	5490113	972996	1126	864.12	-249021
A2:G3	5488354	3155131	1486	2123.24	569320
A3:G1	5389241	775296	1170	662.65	1919
A3:G2	5388275	753293	1138	661.94	-27215
A3:G3	5389058	848719	1287	659.46	-341960
Mixed	5477148	2548294	8319	306.32	370156

tem score is paramount and overrides value for money considerations, in which case it may be useful for emotional characters A1:G1, A2:G1 and A3:G1.

Finally, from table 10, we can see that education strategy 4 offers considerable advantages for the most dysfunctional emotional characters: A1:G2, A1:G3 and A2:G3, and for the mixed population. Again the number of interventions is significantly reduced, although not by as much as education strategy 1, which increases tolerance (as shown in table 7).

From these results we can see that the education strategies that are likely to have a beneficial impact are those which increase tolerance or responsiveness i.e. strategies 1 and 4.

6 DISCUSSION

From tables 7-10 we can derive the summary shown in table 11, which shows for each emotional character the best intervention strategy in terms of total score and gain per intervention. For three emotion characters: A2:G2, A3:G2 and A3:G3, there are no gains from education either in terms of overall score or in terms of gains per intervention. In other cases, increasing tolerance (education strategy 1) tends to offer the best value for money by reducing the number of interventions required, whereas increasing responsiveness (education strategy 4) offers the best system scores since this increases the likelihood of developing mutual co-operation. Of course, it is this very increase in mutual co-operation which sets the scene for a unilateral defection that can result in the need for intervention.

The results presented so far relate to a simulation of 500 rounds. In broad terms the conclusions were not affected when 1000 rounds were played: the sum-

Table 11: Results summary.

Ch.	Best Edu. Str. for Score	Edu. Gain	Best Edu. Str. for Value	Gain per. Int.
A1:G1	3	596	1	762.54
A1:G2	4	1430494	1	868.07
A1:G3	4	1429962	1	868.07
A2:G1	3	250211	-	685.99
A2:G2	-	-	-	390.92
A2:G3	4	569320	1	2546.29
A3:G1	3	569320	-	683.20
A3:G2	-	-	-	675.18
A3:G3	-	-	4	659.46
Mixed	4	370156	1	1399.76

mary given in table 11 also applies to the longer versions. Where there is a significant impact is on the gain per intervention: for homogeneous populations these approximately double as the number of rounds double (increases between 48% for emotional character A2:G2 and 55% for emotional character A3:G3). From this we can say that, the longer the period of concern, the more worthwhile it is to try to improve emotional characters of agents. Increased responsiveness for the mixed population was something of an exception giving an increase in gain per intervention of only 23% compared with the 53% increase in gain per intervention for increased tolerance.

Whilst the ideal from society's point of view is mutual co-operation, this will inevitably lead to intervention so turn-taking, which is maintenance free, may be an acceptable alternative. Consequently, the most cost-effective education strategies may be those that produce emotional characters that form symmetric and asymmetric turn-taking interactions with others. Society's acceptance of turn-taking is likely to increase as simulation length and the likelihood of hopeful defection increase.

7 CONCLUSIONS

What we have attempted to show through these simulations is that the emergence of a successful society, in terms of co-operation amongst its members, is highly dependent on the members' emotional characteristics. Thus while the emotions of anger and gratitude can assist the emergence of co-operation, if the thresholds of activation are not well set, the desired effects do not emerge and the agents become locked in an unbreakable sequence of mutual defection. In such circumstances a society may need to intervene to correct

this detrimental behaviour. The need becomes even more pressing when there are forces working against the development of co-operative relations, here represented by the emotion of hope and the temptation to defect to which it gives rise.

In many cases, however, simply correcting unwanted behaviour is not enough: action is also required to try to prevent recurrence of the unwanted behaviour. Often, there are several ways in which the emotional characteristics of agents can be changed to achieve this; the strategy selected, however, may depend on the aims and emotional characters concerned. Here we have seen that the interventions that provide the best value in terms of gains per intervention do not necessarily lead to the best overall system score. Moreover, the nature of the changes, or even whether changes are worthwhile, will depend on the emotional characteristics of the initial population. In homogeneous populations it will be possible to tailor the education to the population, but in a mixed population this may not be possible.

Of course, what we have reported here applies to a single simulation environment with single parameters for aspects such as the likelihood of defection once hope is activated; changing these parameters may significantly affect the results. However, we believe that, as with increasing the length of the simulations, the effects are likely to be quantitative rather than qualitative. Thus we would expect that the greater the chance of defection once hope is activated, the more essential it will be to change the emotional character of the most dysfunctional agents to keep the number of interventions reasonable; if the probability of defection is small, this becomes less important. Further simulations will be run to confirm these conjectures.

Whilst simulating a society of agents is a good way of investigating the emergence of co-operation, in practice, societies do not allow agents to act as they choose. Societies have a vision for what they deem to be desirable behaviour and will actively intervene to impose upon or persuade their members to adopt - or at least conform to - this desirable behaviour. We believe that it is therefore important that such interventions are themselves modelled, so that the interaction of agents within a society can be better understood.

REFERENCES

- Adam, C., Herzig, A., and Longin, D. (2009). A logical formalization of the OCC theory of emotions. *Synthese*, 168(2):201–248.
- Axelrod, R. (1984). *The Evolution Of Cooperation*. Basic Books, Inc.
- Bazzan, A. L. C. and Bordini, R. H. (2001). A framework for the simulation of agents with emotions: Report on experiments with the Iterated Prisoner's Dilemma. In *Proceedings of the 5th International Conference on Autonomous Agents*.
- Burghouts, G. J., Heylen, D., Poel, M., Akker, R. O. D., and Nijholt, A. (2003). An action selection architecture for an emotional agent. In *Recent Advances in Artificial Intelligence*, pages 293–297. Menlo Park: AAAI Press.
- DeSteno, D., Bartlett, M. Y., Baumann, J., Williams, L. A., and Dickens, L. (2010). Gratitude as moral sentiment: Emotion-guided cooperation in economic exchange. *Emotion*, 10(2):289–293.
- Fehr, E. and Gächter, S. (2002). Altruistic punishment in humans. *Nature*, 415:137–140.
- Hirshleifer, J. (1987). *The Latest on the Best: Essays in Evolution and Optimality*, chapter On the Emotions as Guarantors of Threats and Promises, pages 307–326. MIT Press.
- Lazarus, R. S. (1999). Hope: An emotion and a vital coping resource against despair. *Social Research*, 66(2):653–678.
- Lloyd-Kelly, M. (2014). *Modelling Emotions and Simulating their Effects on Social Interactions in Agent Systems*. PhD thesis, University of Liverpool.
- Lloyd-Kelly, M., Atkinson, K., and Bench-Capon, T. (2012a). Developing co-operation through simulated emotional behaviour. In *13th International Workshop on Multi-Agent Based Simulation*.
- Lloyd-Kelly, M., Atkinson, K., and Bench-Capon, T. (2012b). Emotion as an enabler of co-operation. In *Proceedings of the 4th International Conference on Agents and Artificial Intelligence*, pages 164–169.
- Ortony, A., Clore, G. L., and Collins, A. (1988). *The Cognitive Structure of Emotions*. Cambridge University Press.
- Simpson, B. (2003). Sex, fear, and greed: A social dilemma analysis of gender and cooperation. *Social Forces*, 82(1):35–52.
- Snyder, C. R. (2002). Hope theory: Rainbows in the mind. *Psychological Inquiry*, 13(4):249–275.
- Steunebrink, B. R., Dastani, M., and Meyer, J.-J. C. (2007). A logic of emotions for intelligent agents. In *22nd Conference on Artificial Intelligence*, pages 142–147. AAAI Press.
- Vohs, K. D., Mead, N. L., and Goode, M. R. (2006). The psychological consequences of money. *Science*, 314:1154–1156.