

Stability of Cooperation in Societies of Emotional and Moody Agents

Joe Collenette¹, Katie Atkinson¹, Daan Bloembergen² and Karl Tuyls¹

¹Department of Computer Science, University of Liverpool, UK

²Intelligent and Autonomous Systems, Centrum Wiskunde & Informatica, Amsterdam, NL
j.m.collenette@liverpool.ac.uk

Abstract

It is well documented that cooperation may not be achieved in societies where self-interested agents are engaging in Prisoner's Dilemma scenarios. In this paper we demonstrate, in contrast, that agent societies that use human-inspired emotions within their decision making, can reach stability in cooperation. Our work makes use of the Ortony, Clore, and Collins (OCC) model of emotions and we analyse the evolutionary stability of two different implementations that make use of key emotions from this model. Firstly, we consider an agent society that solely make use of this model of emotions for the agents' decision making. Secondly we look at a model that extends the emotional agents with a model for representing mood. We set out a proof that shows that our emotional agents are an evolutionarily stable strategy when playing against a worst-case scenario strategy. The proof demonstrates that our established model of emotional agents enables evolutionary stability to be achieved, without modification to this model. In contrast, the model of moody agents was shown not to be an evolutionarily stable strategy. Our analysis sheds light on the nature of cooperation within agent societies and the useful role that simulated emotions can play in the agents' decision making and the society as a whole.

Introduction

Models of emotion, particularly those based on the Ortony, Clore, and Collins (OCC) account of emotions, have been used as part of agents' decision making processes to explore their effects on cooperation within social dilemmas (Ortony et al., 1990; Lloyd-Kelly et al., 2014; Collenette et al., 2017). These studies include exploration of the model against a number of different agents within multiple environments. Knowing how these agents, using the OCC model of emotions, react in a simulated environment is an important study in understanding how they will behave against a specific subset of characteristic agents and environments. This paper expands our knowledge of agents using the OCC model of emotions in broader terms, by exploring the evolutionary stability of these agents in a Prisoner's Dilemma setting. The knowledge of whether these strategies can be considered an evolutionarily stable strategy (ESS) allows us to state that these human-inspired agents are able to flourish as a society against invading strategies.

We analyse two different interpretations of OCC agents. Firstly, Emotional agents that decide their action using only the model of emotions, based on a subset of the OCC model. To analyse the evolutionary stability of these agents we use the Prisoner's Dilemma game as this allows us to effectively look at the cooperation these societies achieve and whether the cooperation is sustainable against invading strategies.

We contrast the results with the second interpretation of an OCC agent, the Moody agent (Collenette et al., 2017), which uses the OCC model of emotions for decision making in addition to a psychology-grounded model of mood. Our analysis highlights the different strategies that are needed to achieve success as a society in terms of both stability and cooperation, in the iterated Prisoner's Dilemma.

We find that Emotional agents can be considered an ESS in the Prisoner's Dilemma when given time to converge against the opponent. This is in contrast to Moody agents, which are not an ESS. The Emotional agents take a more defensive strategy that allows cooperation to remain stable over time when playing against other Emotional agents. Opponents are able to take advantage of the Moody agents as they try to create cooperation. Actively trying to create cooperation with opposing strategies makes the Moody strategy more unstable when compared to the Emotional agents.

Analysing two different human-inspired agent strategies, has allowed us to show the inherent risk that cooperation brings in the Prisoner's Dilemma. The work has also showed how different strategies are better deployed in different scenarios. Emotional agents are better suited to a mixed group of agents with differing strategies than the Moody agents, while Moody agents are better suited than Emotional agents when only one strategy exists.

Background

The Prisoner's Dilemma is a social dilemma, popularised through the influential Axelrod's tournament (Axelrod and Hamilton, 1981), where two players pick between cooperating with the other player, or trying to take advantage. This choice is made simultaneously, independent of one another, and with no prior communications. If both players choose

C, C	D, D	D, C	C, D
R, R	P, P	T, S	S, T
3, 3	1, 1	5, 0	0, 5

Table 1: Prisoner’s Dilemma payoff matrix with example payoffs. *Cooperate, Defect.*

to cooperate this is the best joint payoff. If one player defects then the defector receives the best individual payoff, and the cooperator receives the worst individual payoff. If both agents choose to defect they both receive the joint worst payoff. The payoffs for the iterated Prisoner’s Dilemma (Table 1) must have the following restrictions hold (Rapoport et al., 1965) to be valid:

- $(T)emptation > (R)eward > (P)unishment > (S)ucker$
- $R > (S + T)/2$

A strategy can be described as an ESS when the majority of the agents in a society are using this strategy, and it cannot be invaded by any other strategy that is initially rare (Smith and Price, 1973); the definition is given below.

Definition 1. *Let $V(A, B)$ be the expected payoff strategy A receives when playing against strategy B . A strategy M is considered an ESS where M is the dominant strategy and the following holds for all invasion strategies I where $M \neq I$.*

$$V(M, M) > V(I, M) \text{ OR} \\ (V(M, M) = V(I, M) \text{ AND } V(M, I) > V(I, I)) \quad (1)$$

Evolutionary stability in the Prisoner’s Dilemma has been extensively analysed (Thomas, 1985; Bloembergen et al., 2015), with no pure strategy being an ESS in the iterated version of the Prisoner’s Dilemma (Boyd and Lorberbaum, 1987). Furthermore no TIT-FOR- n -TATS is an ESS (Farrell and Ware, 1989), nor are any reactive strategies (Lorberbaum, 1994). The predictability of these kinds of strategies allows invasion strategies to be successful. This shows that evolutionary stability is an extremely demanding criterion to place on a strategy in the iterated Prisoner’s Dilemma. Our Emotional and Moody agents differ from previous analyses of strategies in this problem set, as they are able to identify individual opponents and change their actions based on the individual, and also the memory space of these strategies extends beyond a single interaction.

Emotional agents

We analyse theoretically the evolutionary stability of Emotional agents described by Lloyd-Kelly (2014), whereas Lloyd-Kelly (2014); Collenette et al. (2017) focus primarily on experimental studies only. These emotional agents simulate a subset of the OCC model of emotions (Ortony et al., 1990) from the psychology literature.

The authors of the OCC model show how emotions influence a change in behaviour, with peoples’ actions being a result of their current emotional makeup. A person’s current emotional makeup is then influenced by interactions with other people in the environment. The OCC model defines 22 emotions, which are organised into a hierarchical structure with definitions on how each emotion responds to different actions (Ortony et al., 1990). The OCC model does not define how emotions are processed internally, but gives the outward effects of these emotions. Various agent designers have used this model successfully as part of their agents’ decision making process (André et al., 2000; Popescu et al., 2014; Lloyd-Kelly et al., 2014; Collenette et al., 2017).

Our focus is on two key emotions from the OCC model, *Gratitude* and *Anger*, within agents that interact in the Prisoner’s Dilemmas. The use of these two emotions affects the decision making process of the agents and changes whether an agent is currently cooperating or defecting. Work has also been conducted where additional emotions have been implemented, with the focus being on other areas of agent interactions such as agent replication (Lloyd-Kelly et al., 2014; Lloyd-Kelly, 2014). We are focusing on the decision making strategy of the agents, and whether this can be considered an ESS. This leads us to include only the emotions that affect the outcome of a decision, namely *Gratitude* and *Anger*, this is the same as in Lloyd-Kelly et al. (2014); Collenette et al. (2017).

For clarity we repeat how the two emotions, *Gratitude* and *Anger*, are implemented by Lloyd-Kelly (2014). Each emotion has a value and a threshold, both are represented by an integer. An interaction can increase an emotion’s value, up to the emotional threshold. When the threshold is reached, the agent’s action will change to reflect the emotional trigger, then the value of the emotion is reset. An agent’s threshold is determined by its “personality”. If an Emotional agent receives a cooperative move, the Gratitude value will increase. When the Gratitude threshold is reached the agent will then cooperate with that opponent. Similarly with Anger, if an opponent defects against the Emotional agent, the Anger value will increase. Once the Anger threshold is reached the Emotional agent will now defect against the opponent.

The possible personalities that the previous work has analysed are shown in Table 2. While the thresholds can be defined as much larger, we restrict them to this small set as this best reflects how emotions are short-term biological processes (Keltner and Gross, 1999).

Moody Agents

The Moody agents use simulated model of mood, in addition to the same subset of simulated emotions the Emotional agents use. This mood model informs decision making and is grounded in psychology (Collenette et al., 2016). The mood model uses mood as a real number, lower val-

Anger Threshold	Gratitude Threshold	Character
1	1	Responsive
1	2	Active
1	3	Distrustful
2	1	Accepting
2	2	Impartial
2	3	Non-Accepting
3	1	Trustful
3	2	Passive
3	3	Stubborn

Table 2: Personalities for agents using the OCC model of emotions, with their Anger and Gratitude thresholds

ues represent more depressed moods and higher values represent more positive moods, reflecting how psychologists have represented human mood (Diener et al., 1985; Hepburn and Eysenck, 1989; Hertel, 1999; Bilderbeck et al., 2016). Agents using the mood model use a real number between 0 and 100 to represent mood (Collenette et al., 2016), as this gives an intuitive understanding of whether the current mood is positive or negative. This value can be integrated more accessibly with the other parts of the model.

The model uses the Homo Egualeis concept of fairness (Gintis, 2000) to control how the mood is affected after any given interaction with other agents. Fairness is an important concept when considering whether a given outcome can be considered positive or negative, as fairness has been shown to affect decision making in human societies (Fehr and Schmidt, 1999). Mood has an effect on what is considered fair by affecting perception and judgement (Mayer and Hanson, 1995; Forgas and Bower, 1987). When modelling multi-agent systems, we can capture a notion of human fairness (de Jong et al., 2008) and using this human-inspired model can be beneficial, as de Jong and Tuyls (2011) describes.

Definition 2 shows how the mood value is calculated after an interaction with another agent (Collenette et al., 2016). The value of the mood is affected by the payoff the agent has received and how “fair” the Moody agent believes this payoff to be in the context of both agents’ previous payoffs. Evaluating fairness is achieved using the Homo Egualeis equation (Fehr and Schmidt, 1999). In Equation 2 the first line retrieves the value of α that will be used in the Homo Egualeis equation. By definition this needs to be a value greater than 0 and less than 1, where lower values put less emphasis on the difference between the two agents. $\alpha = \beta$ represents an idealistic scenario where agents care equally about inequity between the opponents and themselves. Higher mood values should give a lower α , as being in a low mood represents that the agent “thinks” it is doing badly in the environment and as such will care more about inequity by design of the mood model (Collenette et al., 2016). This reflects how people also care more about in-

equity when doing poorly (Fehr and Schmidt, 1999). The value of the mood will lie between 0 and 100 and we “flip” the number and divide by 100, for example a mood of 75 will give an α of 0.25.

Definition 2 (Mood Calculation(Collenette et al., 2016)). *Let AG be the set of all agents, with i and $j \in AG$. Let p_i return the payoff of agent i . Let m_i return the mood of agent i , in the range $0 < m < 100$. Let μ_i denote the average payoff for agent i . Let j be the opponent of agent i . Let $\alpha = \beta$.*

$$\begin{aligned}
\alpha_i &= (100 - m_i)/100 \\
\Omega_i(j) &= \mu_i - \alpha_i \cdot \max(\mu_j - \mu_i, 0) - \\
&\quad \beta_i \cdot \max(\mu_i - \mu_j, 0) \\
m_i &\leftarrow m_i + (p_i - \mu_i) + \Omega_i(j)
\end{aligned} \tag{2}$$

The second line is the Homo Egualeis equation. While traditionally the average between the agent and all opponents is taken into account, the model uses only the current opponent so that the agent’s mood is not affected by opponents it never interacts with. The result of Ω is the average payoff the agent has received with a weighted adjustment made based on the difference between the agent’s average payoff and the opponent’s average payoff.

The third and final line shows how the mood will go up or down based on the difference between the payoff the agent received and the agent’s current average payoff, adjusted by the value of the Homo Egualeis equation. The value is then restricted to lie between 0 and 100. This value places the mood into one of five mood levels (Very High, High, Neutral, Low, Very Low).

The mood value will then affect the action selection of the agent, reflecting the psychology literature regarding low moods (Haley and Strickland, 1986; Hertel et al., 2000). For generally low moods, the agent will defect against any new agents. When the mood is high the agent cooperates with new agents, and when the mood is very high, the agent will always cooperate. When the mood is very low, the agent will always defect. The full description of the model along with the psychology grounding is given in Collenette et al. (Collenette et al., 2016, 2017). Table 3 also outlines the changes that Mood causes.

Evolutionary Stability Analysis

To analyse whether Emotional and Moody agents can be considered an ESS, we need to design an opponent strategy that will take the largest advantage of these agents and minimise their payoff. By designing such a strategy we can show that if Emotional and Moody agents are able to remain the dominant strategy, then no other strategy can invade the Emotional or Moody agents.

We will use a strategy termed the *Oracle*. The effectiveness of the strategy is achieved by breaking an assumption of

Mood Level	Moody Agent Cooperating	Moody Agent Defecting
Very High ($m > 90$)	No change	Cooperate
High ($70 > m \geq 90$)	No change	Cooperate against new opponent
Neutral ($30 \geq m \leq 70$)	No change	No change
Low ($10 \leq m < 30$)	Defect against new opponent	No change
Very Low ($m < 10$)	Defect	No change

Table 3: How simulated mood changes the action selection in Moody agents

the Prisoner’s Dilemma, namely that players have no knowledge of the opponent’s move, as reflected by the name. Intuitively the Oracle strategy will always cooperate with itself, and when faced with another strategy will choose the worst outcome for the opponent, effectively making it the worst case scenario for the opponent. The Oracle strategy targets the conditions needed to be an ESS, allowing effective analysis of evolutionary stability.

For example, if an opponent chooses to cooperate, the Oracle strategy is guaranteed to defect, giving the Oracle strategy the T payoff and the opponent the S payoff. For a society of agents to successfully survive an Oracle invasion, that society must have perfect cooperation among themselves, and protect themselves from the Oracle by always defecting against the opposing strategy.

We can now state that the expected value V after one round, for the Oracle strategy o against strategy b , where $Ac(b, o)$ returns the action b would use against o , can be calculated as:

$$V(o, b) = \begin{cases} R & \text{IF } b \equiv o \\ T & \text{IF } b \not\equiv o \text{ AND } Ac(b, o) = C \\ P & \text{IF } b \not\equiv o \text{ AND } Ac(b, o) = D \end{cases}$$

The Oracle is the most effective strategy at minimising the payoff of the Emotional agents, which we show in Theorem 4. To prove this theorem we need to use the fact that Emotional agents will not change their action if their opponent uses the same action (Lemma 3).

Lemma 3. *An Emotional agent will not change its subsequent action against an opponent if its opponent’s action mirrors the Emotional agent’s action.*

Intuitively this means that if an Emotional agent is cooperating and its opponent is also cooperating then the Emotional agent will not switch to defection and visa-versa.

Theorem 4. *The expected payoff of Emotional agents using the defined personalities, in the Prisoner’s Dilemma with the payoffs defined in Table 1, is minimised by the Oracle strategy, with no other strategy being able to lower the expected payoff further.*

Proof. If the Emotional agent is initially defecting then the payoff achieved by the Emotional agent is $V(e, o) = Pn$ where n is the number of rounds. Neither the Oracle nor the Emotional agent will ever change their action, as per Lemma 3.

When the Emotional agent is initially cooperating then the payoff the Emotional agent receives is S as the Oracle defects. By the definition of the Emotional agent we know that the Emotional agent will change to defection when the Anger level of that agent reaches the Anger threshold. The Emotional agent will change its action to defecting against the Oracle agent. Once the Emotional agent has changed its action the Oracle will continue to defect, now both agents are defecting. As both agents are defecting they will continue in mutual defection indefinitely as per Lemma 3, and the Emotional agent will receive the P payoff. We can now state that the expected value of an initially cooperative Emotional agent against the Oracle is $V(e, o) = Sm + P(n - m)$ where m is the Anger threshold and n is the number of rounds.

Assume there is a strategy x where the payoff achieved by the Emotional agent is $V(e, x) < Sm + P(n - m)$ when initially cooperating and $V(e, x) < Pn$ when the Emotional agent is initially defecting.

If the strategy x only defects then the payoff of an initially defecting Emotional agent is $V(e, x) = Pn$, and for an initially cooperative agent $V(e, x) = Sm + P(n - m)$. This contradicts the assumption as $V(e, x) < Sm + P(n - m)$, therefore only defecting is not the strategy x .

If the strategy x only cooperates then the expected payoff of an initially defecting agent is $V(e, x) = Tg + R(n - g)$ where g is the Gratitude threshold. The payoff for the initially cooperative agent is $V(e, x) = Rn$. This leads to a contradiction as T and R are both larger than S and P in the Prisoner’s Dilemma. Strategy x therefore cannot only cooperate.

Strategy x must therefore be a mixed strategy. By Lemma 3 we know that repeating the Emotional agent action leads to an indefinite repetition. Therefore we consider the strategy of doing the opposite of the Emotional agent. When the Emotional agent is cooperating the strategy x will defect and when the Emotional agent is defecting the strategy x will cooperate. Therefore the expected value of the Emotional agent is $V(e, x) = (Tg + Sm) \frac{n}{g+m}$.

This is the minimal strategy since if the strategy x switches a cooperative move for a defect, then the Emotional agent will receive a P payoff, but will not switch to cooperation, effectively removing the S payoffs it would have received. By definition of the Prisoner’s Dilemma $P > S$ so the Emotional agent’s expected value will increase. If the strategy x switches a defection for a cooperative move, then the Emotional agent will receive a R rather than a S payoff, and will not switch to defection leading to further R payoffs. By definition $R > S$ so the Emotional agent’s

expected value will increase.

As strategy x must be the mixed strategy of doing the opposite action of the Emotional agent, therefore the following must hold:

$$V(e, x) < V(e, o)$$

We evaluate this equation where $V(e, x) = (Tg + Sm)\frac{n}{g+m}$ and $V(e, o) = Sm + P(n - m)$ with an initially cooperating Emotional agent. We choose the initially cooperating agent as $Sm + P(n - m) < Pn$.

Given that we are looking at the possible personalities in Table 2 and using the values for the Prisoner's Dilemma in Table 1, we use the personalities that will play the maximum number of cooperative moves and the lowest amount of defection moves to minimize the number of T payoffs and maximize the number of S payoffs. This is the personality Trustful. Plugging in the values in the above equation results in the following:

$$\begin{aligned} (5 \cdot 1 + 0 \cdot 3) \frac{n}{4} &< 0 \cdot 3 + 1(n - 3) \\ \frac{5}{4}n &< n - 3 \\ \frac{1}{4}n &< -3 \\ n &< -12 \end{aligned}$$

For the equation to hold, strategy x must yield a lower expected payoff to the emotional agent than the Oracle, for any number of round n . We have reached a contradiction as n must be positive by definition.

Therefore strategy x is not the mixed strategy. We have also shown that only cooperating and only defecting are not strategy x .

\therefore The expected payoff of Emotional agents using the defined personalities, in the Prisoner's Dilemma with the payoffs defined in Table 1 is minimised by the Oracle strategy, with no other strategy being able to lower the expected payoff further. \square

The Oracle agent is the most effective agent at minimising the expected payoff of the Emotional agents, for the given personalities and the given values for the Prisoner's Dilemma. By restricting the analysis to the given agent personalities and values, we are able to analyse the Emotional agents more effectively as we only need to look at one opposing strategy rather than both an Oracle strategy and the most effective mixed strategy.

Emotional Agents

We now move on to show that the Emotional agents are not an ESS when there are no restrictions on reproduction and interaction speed. We will show this both for the initially cooperative Emotional agent and the initially defecting Emotional agent. The values the strategies will receive for both

the initially cooperative Emotional agent, initially defecting Emotional agent, and the Oracle agent are given in Table 1, for an initial interaction.

Theorem 5. *Emotional agents are not an ESS in the initial phase.*

Proof. Assume that Emotional agents are an ESS. Given a majority of Emotional agents, with an invasion force of Oracle agents, by definition of an ESS, Equation 1 must hold for the Emotional agent strategy M and the invading Oracle strategy I . For the initially cooperative Emotional agent we have,

$$R > T \text{ OR } (R = T \text{ AND } S > R)$$

and for the initially defecting Emotional agent

$$P > P \text{ OR } (P = P \text{ AND } P > R).$$

A contradiction has been reached for each line in both the initially cooperative Emotional agent and the initially defecting Emotional agent. $P > P$ is a contradiction, $T > R$ and $R > P$ (from Table 1) contradict the equations.

\therefore Emotional agents are not an ESS in the initial phase. \square

Emotional agents are not an ESS, due to how the agents respond initially to the Oracle strategy. Emotional agents are able to respond to the opponent on an individual agent level, that is, the action the Emotional agent gives depends on who the opponent is. Collenette et al. (2016) has shown that all initially cooperative Emotional agents will cooperate with each other indefinitely.

Lemma 6. *All Emotional agents will converge to defection with all Oracle agents given a sufficiently high number of interactions and sufficiently high randomness in pairing.*

Emotional agents will eventually choose to defect indefinitely against the Oracle strategy given enough time to adjust. As other agents can not affect the action choice of either the Emotional agent or the Oracle, the Oracle will only defect which only increases the agent's Anger value. The Emotional agent will not switch back to cooperation.

Theorem 7. *Initially cooperative Emotional agents that have fast interactions and slow reproduction are an ESS*

Proof. Assume a fraction ϵ of the population is replaced by the invading Oracle strategy. We also assume that interactions between all agents are fast and reproduction of the population is slow. Given the fast interactions, and slow reproduction with respect to time, all Emotional agents will be defecting against any other Oracle agents that may be residual in the populous, as per Lemma 6. This gives both the Oracle strategy and the Emotional agents the P payoff. No Emotional agent has adjusted to the newly invading ϵ -Oracle agents, and as such are able to receive the S payoff.

Thus the expected payoff of the Emotional agents against the Oracle agents is $V(e, o) = S\epsilon + P(1 - \epsilon)$, and the expected payoff for the Oracle agents against the Emotional agents will therefore be $V(o, e) = T\epsilon + P(1 - \epsilon)$.

Using these values in Equation 1 gives us the following:

$$R > T\epsilon + P(1 - \epsilon) \text{ OR} \\ (R = T\epsilon + P(1 - \epsilon) \text{ AND } S\epsilon + P(1 - \epsilon) > R)$$

The equation will therefore hold, given that ϵ is sufficiently small as per the definition of an ESS strategy (Eshel, 1983). The expected value that the Oracle agent gets from the Emotional agents will be sufficiently close to P such that the first line will always hold. The Emotional agents are protecting themselves from defection of the Oracle agents. The ϵ number of new Oracle agents are unable to take a large enough advantage of the Emotional agents that they can break the stability.

\therefore Initially cooperative Emotional agents are an ESS, when interactions are fast and reproduction is slow. \square

In summary, initially cooperative Emotional agents are an ESS, as no strategy is able to minimise the payoff of the Emotional agents more than the Oracle agent. The assumptions of fast interactions and slow reproduction, are to allow the Emotional agents to adapt to all the Oracle agents before the next reproduction. A sufficiently small epsilon in this case is less than half, if using Table 1 as the payoff matrix, given the assumption that the Emotional agents are the majority as per the definition of an ESS (Eshel, 1983). The assumptions of fast interaction and slow reproduction are part of an efficient evolution and learning process (Hinton and Nowlan, 1987), with fast interactions allowing the agent to learn which in turn guides the reproduction process.

When we consider the initially defecting Emotional agent, they have already adapted to the invading Oracle agents. However the initially defecting Emotional agent being able to protect its payoff is not enough for it to be considered an ESS. The initial defection will prevent the Emotional agents from cooperating as a group, and with no avenue of breaking the defection, this allows the Oracle agents which do work together, to be a fitter strategy.

Moody Agents

We will now be comparing the Emotional agents to Moody agents, again using the Oracle strategy. The Moody agents we are analysing are similar to Emotional agents as they both use OCC-inspired emotions as part of their decision-making process. However the addition of the Mood model on top of the emotions changes how the Moody agents react in certain circumstances. The Moody agents have been shown to perform better in self-play than the Emotional agents (Collette et al., 2017). Analysing the evolutionary stability of these agents shows us the effects the Mood model has, when compared to Emotional agents.

To analyse the Moody agents, we need to take into account that different mood levels affect how the Moody agents respond to the Oracle. Therefore we need to analyse each mood level individually to be able to gain insights into the Mood model as a whole. Table 3 shows each Mood level and when the Mood value will override the action selection of the agent. As we know that Emotional agents need to be initially cooperative to be considered an ESS, we will assume that the Moody agents are also initially cooperative. We will start the proofs from very high levels of mood down to very low moods. For simplicity in the analysis we also assume that the mood levels do not change.

Theorem 8. *Moody agents that are in an initially very high mood are not an ESS*

Proof. Assume Moody agents in a very high mood are an ESS. Given that a fraction ϵ of the population is replaced by the invading Oracle strategy, the expected payoff of two Moody agents in a very high mood is R since by definition all Moody agents are cooperating. Thus the expected payoff of an Oracle agent against the Moody agent will be T and the Moody agent will receive S . Equation 1 holds by definition of an ESS. Therefore the following equation is true:

$$R > T \text{ OR } (R = T \text{ AND } S > R)$$

This is a contradiction, both sides of the *OR* are false. By definition of the Prisoner's Dilemma, $T > R$ which contradicts both sides of the equation.

\therefore Moody agents in an initially very high mood are not an ESS. \square

Moody agents in very high moods are not an ESS. This is due to these particular agents being functionally equivalent to a fully cooperative strategy, which is known to not be an ESS (Boyd and Lorberbaum, 1987; Lorberbaum, 1994). Moving onto high moods, we will show that initially cooperative agents are equivalent to Moody agents in a neutral mood. This allows us to avoid repeating proofs.

Lemma 9. *Moody agents that are in an initially high mood are functionally equivalent to Moody agents in a neutral mood.*

In high moods we know that the only effect on decision making is that the Moody agent will always cooperate with a unknown opponent. By our assumption above, Moody agents are initially cooperative, and therefore Moody agents in a high mood are functionally identical to Moody agents in a neutral mood.

We will show that neutral moods are functionally equivalent to Emotional agents. This gives us that Theorem 7 holds for Emotional agents and Moody agents that are in either a neutral mood or a high mood.

Lemma 10. *Moody agents that are in an initially neutral mood are functionally equivalent to Emotional agents.*

The mood value has no effect on action selection; by definition of Moody agents, they will respond using the emotion model as defined for the Emotional agent. We now go on to show that initially very low, and low, mood levels are not an ESS. We do this as both types of Moody agents are functionally equivalent against an Oracle agent.

Lemma 11. *Moody agents that are in an initially low mood are functionally equivalent to Moody agents against an Oracle agent.*

The Moody agents will defect with all other Moody agents as the initial cooperation is broken. The low moods change the first action to defection, the Moody agents will continue to defect indefinitely as per Lemma 3. The Moody agents in a very low mood will defect by definition. When playing against an Oracle both the Oracle and the Moody agent in a low mood will play defection. The Oracle will also defect against a Moody agent in a very low mood. Regardless of whether the Moody agent is in a low mood or a very low mood, they will defect indefinitely with both other Moody agents and Oracle agents.

Theorem 12. *Moody agents that are in an initially low mood or are in an initially very low mood are not an ESS*

Proof. Assume Moody agents in low mood or a very low mood are an ESS. Given that a fraction ϵ of the population is replaced by the invading Oracle strategy, the expected payoff of two Moody agents is P . This is valid for both very low and low moods as per Lemma 11.

The expected payoff of an Oracle agent against a defecting Moody agent is P , and they will be in mutual defection indefinitely. Equation 1 holds by definition of an ESS. Therefore the following equation is true:

$$P > P \text{ OR } (P = P \text{ AND } P > R)$$

We have reached two contradictions, $P > P$ and $P > R$ since by definition of the Prisoner's Dilemma $R > P$.

\therefore Moody agents that are in an initially low mood or are in an initially very low mood are not an ESS \square

To conclude that Moody agents overall are not an ESS, we need to now show that the Moody agents' mood level will always lead to the evolutionarily unstable mood levels. As only neutral and high moods are possibly an ESS, we only consider these two mood levels.

Theorem 13. *Moody agents in an initially neutral or initially high mood will move to the very high mood level, when there is a sufficiently small ϵ invasion of Oracle agents.*

Proof. The expected payoff of a Moody agent(k) in either a neutral or high mood will be $V(k, o) = S\epsilon + P(1-\epsilon)$ against an Oracle agent and $V(k, k) = R$, as Theorem 7 applies as per Lemma 10, since Moody agents are functionally equivalent to Emotional agents. When a Moody agent receives a payoff its mood level updates. We state that $\Omega_i(j) \approx \mu_i$

when two Moody agents interact. The majority of interactions are between two Moody agents, therefore making their averages (μ_i and μ_j) approximately equal, which means Ω is not changing the perception of the reward in Equation 2.

The final calculation updates the mood $m_i \leftarrow m_i + (p_i - \mu_i) + \Omega_i(j)$. With the majority of the interactions being between two Moody agents we can state that $p_i - \mu_i \approx 0$. We know that $\Omega > 0$, therefore the m_i will increase in the majority of cases indefinitely as the invasion of Oracle agents is ϵ small.

\therefore Moody agents in an initially neutral or initially high mood will move to the very high mood level, given a sufficiently small ϵ invasion of Oracle agents. \square

In conclusion Moody agents are not an ESS. While Moody agents may be an ESS in neutral and high moods, with the same conditions as the Emotional agents, the Moody agents will move into the other mood levels that are not an ESS. If the mood level of Moody agents was to stay stable over time, this would go against the design principles of the model (Collenette et al., 2016). The psychological grounding of the moody model requires that mood levels change over time as per the psychology literature (Mayer and Hanson, 1995; Diener et al., 1985).

Conclusion

We have shown that Emotional agents that use a model of emotions as part of their decision-making can be considered an ESS when they initially cooperate with new partners and are able to adapt to an invading strategy before reproducing. We showed that Moody agents using a simulated model of mood alongside the model of emotions as their decision-making process, are not part of an ESS. This is because some mood levels break the assumption that Moody agents cooperate together and will always protect themselves from invading strategies. We tested these human-inspired agents against an Oracle strategy, which we showed was the most effective at minimising the expected payoff of the Emotional agents and can successfully invade the Moody agents.

Collenette et al. (2017) showed that Emotional agents had stable levels of cooperation, which is reflected by these agents being an ESS. Collenette et al. (2017) also showed that Moody agents increased their cooperation over time, with a higher average payoff than the Emotional agents. The ability to increase their cooperation is what causes the mood model strategy not to be an ESS. This shows us how more protective strategies are able to succeed against invasive strategies, whilst when playing against more reciprocating strategies, a more responsive and risky strategy succeeds. In turn, this reflects the inherent risks and rewards in the Prisoner's Dilemma.

Our work poses further questions, namely whether we can reduce the assumptions needed for the Emotional agents. We will be aiming to find out whether the mood model can

be adapted to become an ESS, by changing the secondary decision making process, which is currently the OCC-based model and is not an inherent part to the mood model. Further alterations to the fairness metric, such as changes to the α and β variables may affect the evolutionary stability, which provides a further avenue of investigation.

This body of work is part of the wider literature that considers simulated emotions, mood, and personalities in agent systems, in terms of how they model these human-inspired concepts and the effects these aspects have on wider agent societies. The majority of the literature concerning these kinds of agents focuses on simulations and observing the effects. We have taken a broader view of human-inspired agents by analysing evolutionary stability in an account that implements both Emotional and Moody agents.

References

- André, E., Klesen, M., Gebhard, P., Allen, S., and Rist, T. (2000). Integrating models of personality and emotions into lifelike characters. In Paiva, A., editor, *Affective Interactions*, volume 1814 of *LNCS*, pages 150–165. Springer.
- Axelrod, R. and Hamilton, W. D. (1981). The evolution of cooperation. *Science*, 211(4489):1390–1396.
- Bilderbeck, A., Reed, Z., McMahon, H., Atkinson, L., Price, J., Geddes, J., Goodwin, G., and Harmer, C. (2016). Associations between mood instability and emotional processing in a large cohort of bipolar patients. *Psychological Medicine*, pages 1–10.
- Bloembergen, D., Tuyls, K., Hennes, D., and Kaisers, M. (2015). Evolutionary dynamics of multi-agent learning: a survey. *Journal of Artificial Intelligence Research*, 53:659–697.
- Boyd, R. and Lorberbaum, J. P. (1987). No pure strategy is evolutionarily stable in the repeated prisoner’s dilemma game. *Nature*, 327:58 EP –.
- Collenette, J., Atkinson, K., Bloembergen, D., and Tuyls, K. (2016). Modelling mood in co-operative emotional agents. In *Proceedings of Distributed Autonomous Robotic Systems*.
- Collenette, J., Atkinson, K., Bloembergen, D., and Tuyls, K. (2017). Environmental effects on simulated emotional and moody agents. *The Knowledge Engineering Review*, 32.
- de Jong, S. and Tuyls, K. (2011). Human-inspired computational fairness. *Autonomous Agents and Multi-Agent Systems*, 22(1):103–126.
- de Jong, S., Tuyls, K., and Verbeeck, K. (2008). Fairness in multi-agent systems. *The Knowledge Engineering Review*, 23(02):153–180.
- Diener, E., Larsen, R. J., Levine, S., and Emmons, R. A. (1985). Intensity and frequency: dimensions underlying positive and negative affect. *Journal of personality and social psychology*, 48(5):1253.
- Eshel, I. (1983). Evolutionary and continuous stability. *Journal of theoretical Biology*, 103(1):99–111.
- Farrell, J. and Ware, R. (1989). Evolutionary stability in the repeated prisoner’s dilemma. *Theoretical Population Biology*, 36(2):161 – 166.
- Fehr, E. and Schmidt, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly journal of Economics*, 114:817–868.
- Forgas, J. P. and Bower, G. H. (1987). Mood effects on person-perception judgments. *Journal of personality and social psychology*, 53(1):53.
- Gintis, H. (2000). *Game theory evolving: A problem-centered introduction to modeling strategic behavior*. Princeton university press.
- Haley, W. E. and Strickland, B. R. (1986). Interpersonal betrayal and cooperation: Effects on self-evaluation in depression. *Journal of Personality and Social Psychology*, 50(2):386.
- Hepburn, L. and Eysenck, M. W. (1989). Personality, average mood and mood variability. *Personality and Individual Differences*, 10(9):975–983.
- Hertel, G. (1999). Mood effects in social dilemmas: What we know so far. *Resolving social dilemmas: Dynamic, structural, and intergroup aspects*, pages 227–243.
- Hertel, G., Neuhof, J., Theuer, T., and Kerr, N. L. (2000). Mood effects on cooperation in small groups: Does positive mood simply lead to more cooperation? *Cognition & emotion*, 14(4):441–472.
- Hinton, G. and Nowlan, S. (1987). How learning can guide evolution. *Complex Systems*, 1:495–502.
- Keltner, D. and Gross, J. J. (1999). Functional accounts of emotions. *Cognition & Emotion*, 13(5):467–480.
- 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. (2014). Fostering co-operative behaviour through social intervention. In *Proc of SIMULTECH’14*, pages 578–585. IEEE.
- Lorberbaum, J. (1994). No strategy is evolutionarily stable in the repeated prisoner’s dilemma. *Journal of Theoretical Biology*, 168(2):117 – 130.
- Mayer, J. D. and Hanson, E. (1995). Mood-congruent judgment over time. *Personality and Social Psychology Bulletin*, 21:237–237.
- Ortony, A., Clore, G. L., and Collins, A. (1990). *The cognitive structure of emotions*. Cambridge university press.
- Popescu, A., Broekens, J., and van Someren, M. (2014). Gamygdala: An emotion engine for games. *IEEE Transactions on Affective Computing*, 5(1):32–44.
- Rapoport, A., Chammah, A. M., and Orwant, C. J. (1965). *Prisoner’s dilemma: A study in conflict and cooperation*, volume 165. University of Michigan press.
- Smith, J. M. and Price, G. R. (1973). The logic of animal conflict. *Nature*, 246(5427):15.
- Thomas, B. (1985). On evolutionarily stable sets. *Journal of Mathematical Biology*, 22(1):105–115.