Emergence of social dynamics among affectively motivated agents

Theije Visser

Master's thesis

Artificial Intelligence
Department of Artificial Intelligence
University of Groningen
Groningen, The Netherlands

Supervisors:

dr. T. Andringa
prof. dr. L.C. Verbrugge

January 31, 2011

Abstract

Social simulation research on the emergence of social structure among individual organisms is generally studied by computational models in which agents make rational decisions or no decisions at all. In this work, we develop an agent-based model which simulates emotion as the motivation for decision-making of individual affective agents within a simulated social environment. We formalize an alternative theory of agency by Di Paolo, which states that agents are fully coupled to the environment via perception and action through core affect, a dimensional model of emotion. We propose that emotion lies at the core of the affective self-regulation of the coupling between the agent and the environment. As a first experiment, we use our model to generate the emergence of zones of cooperation in the Demographic Prisoner's Dilemma (DPD) as proposed by Epstein. We hypothesize that our model can provide a cognitively plausible explanation for the dynamics of cooperation within the DPD. In the second experiment, we use our model to generate a sociological phenomenon, called the histeroidal cycle. The theory backing this phenomenon states that, over the course of generations, the influence of a sociopathic minority on the well-being of the cooperative majority of the society increases and decreases periodically. Our model assumptions are sufficient to generate these dynamics and can be applied to gain new insights into the affect-based behavioral foundations of the evolution of social structure in general.

Contents

1 Introduction **1**
1.1 Observations and research questions 2
1.2 Social simulation 3
1.3 Methodology 3
1.4 Structure of the thesis 4

2 Theoretical review **5**
2.1 Target phenomenon 5
2.1.1 Definition of sociopathy 6
2.1.2 The histeroidal cycle 7
2.2 Related research 9
2.2.1 Game-theoretic approaches 9
2.2.2 Demographic Games 12
2.2.3 Agent-based models 13
2.2.4 Agency 14
2.3 Indivaidual agent behavior 17
2.3.1 Core affect 17
2.3.2 Emotion and agency 18
2.3.3 Emotion in agent-based models 20

3 Model design **21**
3.1 Agent design 21
3.2 Environment 23

4 Baseline model **26**
4.1 Agent components 27
4.1.1 Action 27
4.1.2 Conscience 27
4.1.3 Received utility 29
4.1.4 Viability 30
4.2 Verification 30

5 First extension: the perceiving agent **34**
5.1 Agent components 35
5.1.1 Perceiving the environment 35
5.1.2 Action 35
5.2 Verification 36

6 Second extension: The affective learner agent **39**
6.1 Agent components 40
6.1.1 Social skill 40
6.1.2 Skill update 40
6.1.3 Expected utility 41
6.1.4 Arousal 42
6.1.5 Received utility 42
6.2 Verification 43

7 Validation experiments **47**
7.1 Experiment 1: baseline model and extensions without evolution . 48
7.1.1 Baseline agent 49
7.1.2 Perceiving agent 53
7.1.3 Affective learner agent 56
7.2 Experiment 2: baseline model and extensions with evolution . . . 61
7.2.1 Baseline agent 61
7.2.2 Perceiving agent 63
7.2.3 Affective learner agent 64

8 Discussion **68**
8.1 Shortcomings and future work 69
8.2 Conclusion 70

A Graphical user interface **76**

Chapter 1

Introduction

The pursuit of happiness can be seen as the main driving force people base their choices in life on. Our individual preferences and attitudes are, for a great part, caused by the assessment of our current feeling of well-being caused by our experiences. However, the pursuit of happiness of individuals within a society does not necessarily entail the increase of happiness measured over the whole society. A possible cause could be the inability of the co-operating majority of the society to detect and counter-act the minority of sociopathic individuals, which exploit the altruistic tendencies of this majority.

This research aims to investigate the interplay between internal state and environment emerging from affectively motivated behavior among a society of interacting agents. The tool we use is an agent-based modeling framework simulating a heterogeneous population of different types of affective cognitive individuals. The behavior of agents is based on the interaction dynamics between the affective states of the agent and the states of the agents in the environment. Here, we propose that affect regulation is the only motive for behavioral choices in organisms.

We will develop a computational agent-based model which models agents as affect-regulating beings. Agents interact by playing a fixed strategy in games against neighbors, where utilities are based on the Prisoner's Dilemma. Affective states are modeled by *core affect*, from the dimensional theory of emotion as proposed by Russell [38]. Core affect is defined as the simplest non-reflective feeling evident in moods and emotions. It consists of two dimensions, *pleasure* and *arousal*. Affective states are interpreted as a continuous assessment of the current state of oneself. It is thus reasoned that living beings try to optimize the feeling of well-being as it is a direct indicator of the being's state of viability; the distance from death [12].

In this research, our main interest is to explain the emergence of social structure from the micro-behavior of interacting agents which are motivated by the most basic demands of agency and co-operation. Affectively motivated social behavior implements this theory of agency. The Demographic Prisoner's Dilemma paradigm [14] is used as a basis and reference to investigate the dynamics of a society which consists of affective agents.

We will try to simulate the dynamics arising from the interactions between

individuals in a society which consists of different types of agents, according to whether they are biased to egoistic or altruistic strategies [26]. We aim to investigate how the effects of interaction between different types of individuals can be detrimental to the average well-being of the whole society. Specifically we will investigate when the altruistic population is able to counterbalance the negative influence of sociopathic agents by affective self-regulation.

1.1 Observations and research questions

This research originally stems from a number of observations of human society. The first *behavioral* observation is that humans are not fully rational and motivate their decisions based on the expected affective impact of their actions [17]. The second *phenomenological* observation is based on a theory of agency, which states that the feeling of well-being can be linked to the viability of an individual. The third *sociological* observation is that, through periods of history, societies were characterized by oppressive regimes in which few people experience a high level of well-being [26].

The latter observation is the case study of the model. We aim to generate the dynamics of the histeroidal cycle as described by Lobaczewski [26]. He attempts to uncover the sociological dynamics of the rise of power of egoistic individuals by his first-hand observations of the behavior of communists in post-war Poland. We use the term *sociopath* to refer to these evil-doing individuals. More precisely, by sociopaths we mean people with an antisocial/dis-social personality disorder, as diagnosed by the Hare Psychopathy Checklist [19]. We will refrain from using the term psychopath, as this term has grown to be too ambiguous through popular culture. The term sociopath refers to an individual who exhibits behavior that is detrimental to the well-being of people it is interacting with. The observations as mentioned above are caused by a great number of factors. However, in this research, we constrain ourself by only focusing on answering the following general question.

> *How can an affect-based model of human cognition improve the understanding of the motivation of human behavior within a social environment?*

More specifically, we are interested in so-called *pathological* societies. A pathological society is characterized by a minority of sociopathic individuals which causes the majority of the individuals in the society, the cooperating individuals, to feel unhappy. The following sub-hypotheses emphasize these subjects.

1. How does social structure emerge from emotionally motivated behavior among a society of interacting agents?

2. How is the origin of the onset and offset of pathological societies over time explained by the interactions between the individuals of the society?

1.2 Social simulation

Social simulation by computational modeling is one of the most widespread tools to study the foundations of sociological phenomena [18, 15]. Different methods exist within this field, from system dynamics modeling and game theory to cellular automata and agent-based modeling. We draw from the latter approach to model each member of our artificial society separately. Another reason to use the agent-based modeling approach is that we aim to model both the internal affect development of agents and their interaction networks as emergent features of the simulation [40]. By explicitly locating all agents on a two-dimensional grid, the interaction network of an agent is determined by this environment.

To explain the macro-social phenomena as described above, our agent-based model needs to be built as an abstraction of theoretical assumptions of the micro-rules underlying these phenomena. Our model is based on psychological theories about human behavior. For our model, the most important theory is a dimensional theory of emotion called core affect. It represents the two dimensions of pleasure (an assessment of the current condition) and arousal (the energy expenditure available for the current situation).

Decision making is modeled as a function of the predicted core affect [38], by assuming that humans always seek behavioral options that maximize the feeling of well-being. The evaluation of the current core affect and the prediction of future core affect serve as motivation for making choices about which action to perform. The social skill level determines the ability of an agent to effectively predict its core affect in the next state.

Actions are modeled by the movement of the agent on the grid. Interactions between agents are modeled as two-player games with all other agents in the current interaction range. The action choice of an agent depends on the evaluation of the internal state (core affect) and the affective influence of other agents within the interaction range. The interaction range is dynamic, as it depends on the number of agents in the neighborhood and the social space of all agents. The size of the sphere of influence depends on the internal state of the agent.

1.3 Methodology

We develop our agent-based model in steps. This means that a baseline model consists of an implementation of only the most basic features of the theoretical assumptions. Subsequently, the behavioral features of the model are tested by verification experiments. Following these experiments, we can optimize the model so that agent behavior will adhere to our assumptions about affectively motivated behavior. Thereafter the model is extended with additional refinements. Each increment is tested and compared with the baseline model. The validity of the model can now be substantiated in an early stage, and it is easier to extend or modify the final model in future research.

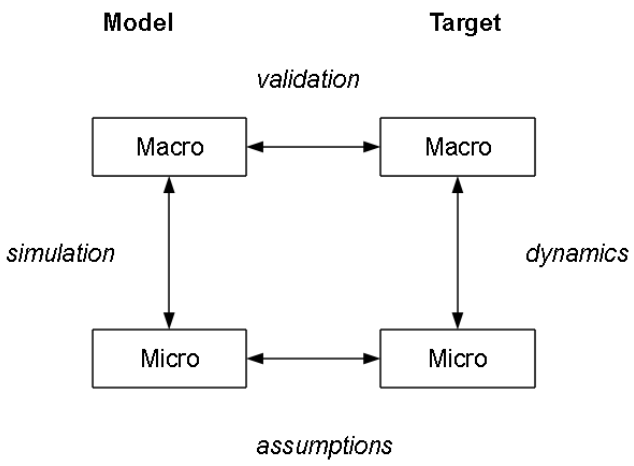

Figure 1.1: General methodology of a social simulation research program.

The final model is intended to explain the above-mentioned macro-social observations by evaluating the output of the model (see figure 1.1). To achieve this, we will first try to replicate the outcomes of the research by Epstein [14], in which agents interact on a torus and receive payoffs according to the Demographic Prisoner's Dilemma (DPD) payoff matrix. If the macro-structure of the model qualitatively resembles the macro-structure of the data generated by the DPD, then the macro-structure can be explained by our theory as implemented in the micro-structure of the model.

The final goal of our model is the generation of the histeroidal cycle as defined by Lobaczewski [26]. The theory of the histeroidal cycle states that the society of agents experiences a periodic movement of internal agent variables over multiple generations. We will investigate whether our society of affectively motivated agents shows this periodic cycle. Conclusions can then be drawn on the basis of our cognitively plausible assumptions about agent behavior.

1.4 Structure of the thesis

Chapter 2 starts with a theoretical review of the psychological foundations of our research. The second part of the chapter focuses on similar research in the field of agent-based modeling aimed at sociological phenomena. Chapter 3 combines the elements discussed in chapter 2 and defines the methodology of the model design and testing procedures. Chapters 4, 5, and 6 describe the formalization of individual agent behavior by increasing levels of complexity. Chapter 7 describes the experimental results of different agent society configurations. Finally, chapter 8 discusses the results of these experiments. The chapter ends with a discussion of future work and a general conclusion.

Chapter 2

Theoretical review

In this chapter we will review relevant research which serves as the theoretical basis for the development of our computational model. We will use theoretical research questions to structure the chapter. The first section will describe the target phenomenon, in other words, the sociological phenomenon which is to be generated by our model.

1. *What are the features of a histeroidal cycle, and which findings point to the existence of this phenomenon?*

The next sections will focus on similar research in the field of social simulation. The following two theoretical research questions constrain our domain of interest:

2. *What are the current methods to simulate the interaction among individuals in a spatial society spanning multiple generations?*
3. *Which assumptions and abstractions do these methods make about the real world phenomena under investigation, i.e. how cognitively plausible are these methods?*

The last sections of the chapter describe our assumptions about individual behavior of social agents. We will discuss the following question.

4. *Which findings from psychological and sociological research can be incorporated in a formal model of the social agent to make its behavior more cognitively plausible?*

The concepts discussed in these sections will form the basis of the formal definition of the agents in our framework. The next chapter will then proceed to describe the actual formalization of these concepts in the individual agent model.

2.1 Target phenomenon

This section will discuss the so-called 'target phenomena', which are the social phenomena we aim to simulate with our agent-based model. Some phenomena

are based on theories, whereas other phenomena are based on empirical data. Our target phenomenon is varying influence of sociopathic individuals on the majority of a cooperative society. Lobaczewski [26] proposes that this varying influence is of a periodic nature. He labels these social dynamics as the *histeroidal cycle*. Before describing the theory behind this phenomenon, this section will now first turn to the nature of egoistic individuals, which we will call *sociopaths*.

2.1.1 Definition of sociopathy

The science of *ponerology*, proposed by Lobaczewski [26], is the study of the root causes and genesis of evil, on both social and interpersonal levels. Lobaczewski identifies various types of pathological individuals. *Schizophrenic* individuals often provide the naive and misguided ideology, *paranoids* are the first to gain leadership positions in ponerogenic groups, and *sociopaths* are the eventual inspirational source for the entire pathocratic system, occupying all positions of influence. Stout [44] points out that it is actions and not motivations that truly count.

When trying to provide a more specific behavioral definition of sociopathy, throughout psychiatric and psychological literature, we see that the meaning of sociopathy is very dispersed (see Reimer [36] for a critique on explicitly defining sociopathy). However, the general consensus is that sociopaths can be defined as individuals exhibiting a diminished *conscience* level [43, 20]. According to Stilwell [42], conscience is defined as a fixed, cultural bias towards cooperative behavior. This is comparable to the definition of Frank of a 'defecting' individual: "A cooperator is someone who, possibly through intensive cultural conditioning, has developed a heritable capacity to experience moral sentiment that predisposes him to cooperate. A defector is someone who either lacks this capacity or has failed to develop it" [16].

Consequently, conscience is the tendency to refrain from making decisions which decrease the viability of the surrounding agents. Conscience can also be correlated with empathy [9]; the higher the conscience, the higher the empathy-level of an agent. Furthermore, low conscience is associated with abusive behavior, while high conscience is associated with cooperative behavior.

The Psychopathy Checklist: Screening Version [19] is a psycho-diagnostic test used to assess the level of psychopathy in individuals from both forensic and civic populations. It consists of 20 items, of which each of the items in the PCL-SV is scored on a three-point scale. It is commonly used for assessing the probability that of rehabilitation of individuals after committing a criminal act.

Coid et al. [7] study the prevalence of psychopathic traits in the household population of Great Britain. In this study, the standard PCL: SV test (Psychopathy Checklist: Screening Version [19]) is used to quantify the prevalence of sociopathy in subjects. Figure 2.1 shows the results of this experiment. The highest possible score is 24.

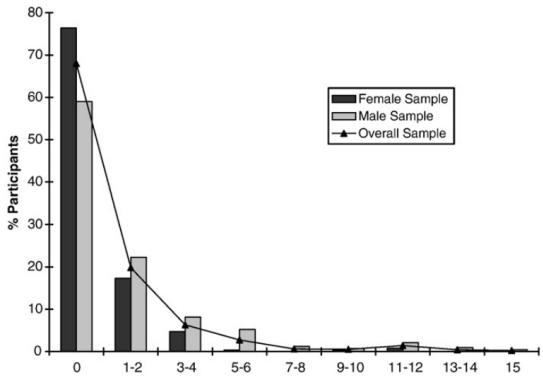

Figure 2.1: Distribution of weighted PCL:SV scores [7]

The figure suggests a 'half-normal' distribution of psychopathic traits. The weighted prevalence of possible sociopathy, using a cut score of 11 or more points in this population was 2.3 percent. Using a cut score of 13 or more points, the weighted prevalence was 0.6 percent. The study by Coid et al. is one of the first studies to provide insight into the proportion of psychopathy in a civic population.

2.1.2 The histeroidal cycle

Now, the following question remains: *What are the behavioral mechanisms which result in a cooperative society to be abused and controlled by this small group of sociopaths?*. In other words, we would like to investigate which kind of micro-behavior causes the oscillating macro-dynamics of the well-being of a society, designated as the histeroidal cycle.

Stout proposes that "[...] the limbic system plays a dominant role in regulating our feelings, the accessibility of our memories, our motivations to act, our ability to make meaning of our experiences, and even our consciences" ([44], p. 77). She defines conscience as a compelling feeling of obligation that is always based in the tendency to bond and cooperate with others. In this way, moral character, or conscience, is causally linked to the capacity to form emotional attachments.

All cooperative individuals are affected by the emotional state of those around us; we can all become traumatized when a small part of society experiences a traumatic event. Sociopaths are the exception; since they are not receptive to emotions, they can live in environments in which people suffer and have fear, without experiencing negative emotional effects. Sociopaths can be seen as people without *empathy*.

"In an abusive relationship, the victim, paralyzed by constant fear, clings to the 'protection' of the very person who terrorizes them" ([44], p.132). An individual is abused by 1) its predisposition to fear and 2) its conviction that because the world is so fearful, loyalty to a protector is necessary. In reality however, the unhealthy relationship is fearful, and the protector is the abuser.

By this 'limbic warfare', large-scale social changes can be initiated by a small group of sociopaths who tap into the anger and paranoia of a vulnerable population. At these points in social history, countries can be ruled by a *pathocracy*, a macro-social disease that can last for generations.

When sociopaths are exposed, and their nature is understood, they are not able to con cooperative individuals anymore. When we lose the ability to recognize pathological behavior, sociopathic individuals are able to influence the majority of society. Everything follows from this inability to accurately read objective reality. The stages of a histeroidal cycle [26] (graphically represented in figure 2.2) are defined as follows.

1. *Group trauma*: a national catastrophe will instate group fear on the cooperating majority.
2. *Spreading of fear*: sociopaths will attempt to use and amplify the already present fear in individuals by starting abusive relationships with frightened individuals.
3. *Revolution*: people become aware of the true nature of sociopaths and the influence of sociopaths is reduced to a minimum.

In our research we are interested in the explanation of the behavioral foundations of the histeroidal cycle (see figure 2.2). We aim to create a model which simulates the histeroidal cycle (qualitatively). We will study the behavior of the agents under different conscience distributions. We will observe under which proportions sociopathic agents cause the most damage to the whole society. Generally speaking, we are interested in the conditions under which sociopaths are given the freedom to control a society.

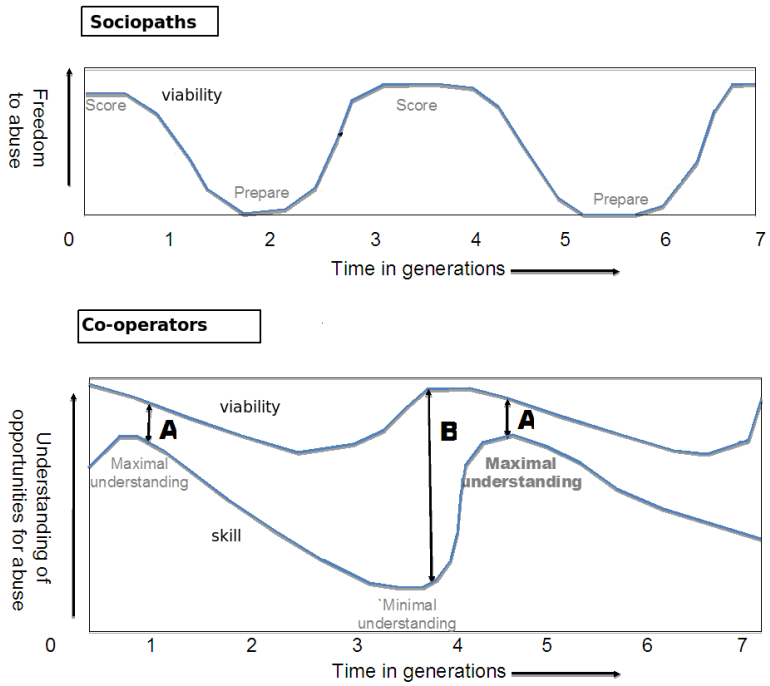

Figure 2.2: The histeroidal cycle as proposed by [26].

The next section will discuss the computational approaches to the study of social dynamics. We will describe established social simulation models which are similar to our approach. After this, we will discuss models of emotion which will form the basis of our formalization of individual agent behavior.

2.2 Related research

Computational models of interactions among different types of individuals within a society is the subject of a great number of publications in the field of social science, theoretical biology, and economics. In this section, we will focus on two classes of computational models aimed at the explanation of social phenomena. First, we will describe the class of game-theoretic models. Second, current agent-based methods for modeling social phenomena are discussed.

2.2.1 Game-theoretic approaches

Game theory, invented in 1944 by Von Neumann and Morgenstern [48], was designed to study human behavior in strategic and economic decisions. Since its inception, it has also been used on a wide range of topics in biology and ecology. In more general terms, game theory studies the ways in which strategic interactions among *rational* agents produce outcomes with respect to the preferences of those agents.

Preferences of agents towards an event are represented by the *utility* function. The utility function designates the change in subjective welfare (*payoff*) that an agent derives from an event. An agent is rewarded or punished each time according to the event.

From now, we will assume that an event is a strategic interaction with one other agent. This kind of interaction is called a *two-person game*. Each player uses a strategy against the other player. The strategy tells the agent what *actions* to take in response to every possible strategy other players might use.

As mentioned, in game theory, an agent is assumed to be economically rational. Firstly, this means that the agent can order outcomes with respect to their contributions to its welfare. Secondly, it can see which sequences of actions would lead to which outcomes. And finally, the agent chooses which action to take according to which action yields outcomes associated with the maximum utility, given the actions of the other players [37]. Thus, the agent makes decisions motivated by reasoning about what seems best for its purposes.

Payoffs in a two-player game can be represented by a *payoff matrix* of dimension 2. Equation 2.2.1 shows such a payoff matrix for the strategies p and q.

	p	q
p	A, A	B, C
q	C, B	D, D

(2.1)

In this matrix, the strategy options (p, q) of the first player are denoted by the left-most column, while the strategy options (p, q) of the second player are denoted by the upper row. Each cell of the matrix gives the payoffs to both players for each combination of actions. The payoff for the first player appears as the first number in the cell; the payoff for the second player appears as the second number. When, for example, both players choose strategy p, then each get a payoff A. When the first player chooses p, and the other player chooses q, the first player receives payoff B, while the second player receives payoff C, and so forth.

From the payoff matrix, each player can evaluate his or her two possible actions by comparing their personal payoffs in each column, since the matrix shows the preferable action, for each possible action by their opponent. Hence, a player's *best action* depends on expectations about the actions of all the other players.

Nowadays, the main business of game theory lies in finding the solution of n-player games. Following the general practice in economics, game theorists refer to these solutions as equilibria. The Nash equilibrium [30] is of importance in the practice of finding analytical solutions to these games. A strategy pair (x, y) if and only if both players cannot deviate from and increase their payoff, given that the other player keeps playing the previous strategy. The concept of the Nash equilibrium meant an important change in the game-theoretic landscape. While Von Neumann and colleagues where concentrating on cooperative games in which players try to maximize the cumulative payoff of all players, Nash was

now looking at *non-cooperative* games, in which player A tries to maximize their own payoff while minimizing B's payoff.

Now that we described the building blocks of game theory relevant for our research, the next section will go into one particular example of an n-person game: the *Prisoner's Dilemma*, which has famously been used to study the origins of cooperation among members of societies.

Prisoner's Dilemma

The problem of cooperation within a society of agents is studied in the Prisoner's Dilemma game [1]. The generalized payoff matrix of for this game is as follows.

	$cooperate$	$defect$
$cooperate$	R, R	S, T
$defect$	T, S	P, P

(2.2)

Here, T stands for Temptation to defect, R for Reward for mutual cooperation, P for Punishment for mutual defection and S for Sucker's payoff. The following condition *has* to hold for a game to count as a Prisoner's Dilemma: $S < P < R < T$. The Prisoner's Dilemma demonstrates why two people might not cooperate even if it is in both their best interests to do so.

In this classic form of the game, the only possible Nash equilibrium for the game is for all agents to defect. Given a fixed strategy of the opponent, an agent will always gain a greater payoff by playing defect. This entails that in a society consisting of only rational players, everyone will defect. The question remains why cooperation is able to persist in real-world societies. We will come back to this question later in this chapter.

Evolutionary game theory

The relation between the dynamics of social strategies and the relative frequencies of these strategies within a society of agents spanning *multiple generations* is the subject of evolutionary game theory. This theory was first proposed by James Maynard Smith [29] in 1982.

Individuals have *fixed* strategies (a phenotype), and they interact randomly with other individuals in the society. Agents replicate depending on the payoff of the strategy they use (fitness). Consequently, success in the game is translated into reproductive success. One time-step corresponds to one generation of agents. These *replicator dynamics* [29, 49, 21] can be modeled by differential equations. The most important equation is called the replicator equation:

$$\dot{x}_i = x_i(f_i - \phi) \tag{2.3}$$

This replicator equation defines the change $\dot{x}_i$ of the relative frequency of a strategy i as a function of the expected payoff f_i, the average payoff ϕ and the current relative frequency x_i. The replicator dynamics are the cornerstone of evolutionary game theory.

To analyze the behavior of a population in an n-player game under selection dynamics without mutation, the following question could be asked. Could a new strategy, which is not present in the population, invade the population and increase in frequency? John Maynard Smith defined the concept of an Evolutionarily Stable Strategy (ESS) to answer this question. Maynard Smith defines ESS as follows:

> (...) an ESS is a strategy such that, if most of the members of a population adopt it, there is no 'mutant' strategy that would give higher reproductive fitness [29].

The ESS concept and the Nash equilibrium are related concepts. We will not go into further mathematical details of the ESS. However, we can state that the most important feature of the replicator dynamics is that their ESSs can correspond to the strategies that would be adopted by fully informed rational players of the game.

For the Prisoner's Dilemma, replicator dynamics lead to pure defection[1]. Even the slightest perturbation from cooperation ultimately results in pure defection [14].

Spatial models

Game-theoretic models can be extended to include spatially-determined interactions. In spatial models, the interaction neighborhood is much smaller than the population as a whole. Agents occupy a cell on a spatial representation of the world. This representation can be a grid with a rectangular boundary, a circle, or surface of a sphere or torus. On this representation, local interactions occur between agents. The deterministic cellular automaton [47, 24] is a prime example of a spatial model. In this model, agents are not able to move, and every time-step they update their 'state' by adapting their strategy to match the most 'successful' strategy (gaining the most payoff) in the neighborhood.

For evolutionary game theory, spatial effects can influence the outcome of frequency-dependent selection. Strategies which would exclude each other in a homogeneous framework can now coexist. The addition of space also introduces new outcomes in the case of the Prisoner's Dilemma. Nowak and May [31, 32] showed that cooperators and defectors *can* coexist in the spatial Prisoner's Dilemma. The main explanation is that cooperators are able to survive in clusters, so that defectors do not have the opportunity to negatively affect cooperators within the clusters.

2.2.2 Demographic Games

Epstein introduces *Demographic Games* [14]. This class of Demographic Games differs from the 'traditional' spatial evolutionary games in a number of ways. First, the agents are not fixed on the grid; each time step they *move* randomly on the grid. Second, the agents maintain an *internal state*. Epstein applies this new class to the Prisoner's Dilemma. He defines his Demographic Prisoner's

[1] All players in the society play the defection strategy.

Dilemma (DPD) model as follows. The agents are randomly distributed on a 30-by-30 torus. The fixed agent strategy (cooperate or defect) is also assigned randomly to each agent.

Every time-step, agents play a game with each neighbor within the Von Neumann neighborhood. From every interaction, the agents receive a payoff from the payoff matrix of the Prisoner's Dilemma (see table 2.2.1. Negative payoffs are used [35]; specifically the following condition has to hold:

$$T > R > 0 > P > S \tag{2.4}$$

The internal state is represented by the accumulation c of payoffs over time. Each step, the agent evaluates c. The agent *dies* when $c < 0$, the agent procreates if $c > t$, where $t > 0$ is the reproduction threshold. The offspring receives a part p of the accumulated c of its parent. The parent's new accumulated payoff is $c - p$. The offspring then appears randomly at one of the cells in the neighborhood. If all cells are occupied, no offspring appears. Agents are all born with a randomly determined age between 0 and 100. Agents die when the age is above 100.

The important difference between the DPD and the cellular automaton implementation of the Prisoner's Dilemma is that the standard replicator dynamic assumes that the frequency of a strategy grows according to the fitness of the strategy relative to the average fitness (see equation 2.3).

Epstein's [14] results show that cooperation emerges in such a population of agents exhibiting fixed strategies and internal states. Epstein also shows that space and local interactions are the cause of this phenomenon, as the system runs to pure defection, consistent with the replicator dynamics of the Prisoner's Dilemma.

2.2.3 Agent-based models

The class of Demographic games is a very simple example of an agent-based model. We will now go into more detail regarding the definition of agents and agent-based modeling. As mentioned earlier, agent-based models are used in the field of social science to explain the emergence of macroscopic societal regularities. The main difference between agent-based models and the aforementioned computational approaches is that they allow each agent to have a unique identity, and they do not require that agents are evenly distributed across space.

There are different approaches to the modeling of artificial societies. The computational agent-based approach models agents by adhering to the principles of heterogeneity, autonomy, explicit space, local interactions and bounded rationality [15]. A Belief-Desire-Intention (BDI) agent [4] implements the idea of bounded rationality. A boundedly rational agent resolves the problem of the need for unlimited computational power when considering rational agents in a multi-agent environment. A simple version of the BDI agent suitable for agent-based modeling makes decisions on the basis of its desires and in reaction to its environment.

The BDI-agent is an example of an intentional system. Daniel Dennett [10] defines this as a system which experiences states *about* other states, things, or events. In our research, agent behavior is implemented from the *intentional*

stance, i.e. we define agent components in terms of goals, intent and feeling. In the next section, we will go into more detail as to the nature of agents and their behavior, by providing a definition of *agency*.

2.2.4 Agency

All living beings exist in precarious environments. Behavior of living beings is therefore in the first place aimed at maintaining existence in a way that keeps the individual (far) from death. Di Paolo [12] calls the distance from death *viability*. Viability is essential for all living entities. An explicit definition of *biological* agency was recently proposed [12, 2]. In this definition, an agent is defined as a self-constituting system that adaptively regulates its coupling with its environment and contributes to sustaining itself as a consequence (see figure 2.3). An agent acts by modulating the coupling between agent state and the environment state. In other words, an agent is able to *adapt* its relation to a precarious environment. Adaptivity is defined on an abstraction level that does not require a neural system to meet the listed functional requirements:

> "Adaptivity is a system's capacity, in some circumstances, to regulate its states and its relation to the environment with the result that, if the states are sufficiently close to the limits of its viability, 1) tendencies are distinguished and acted upon depending on whether the states will approach or recede from these proximal limits and, as a consequence, 2) tendencies that approach these limits are moved closer to, or transformed into, tendencies that do not approach them and so future states are prevented from reaching these limits with an outward velocity."[2].

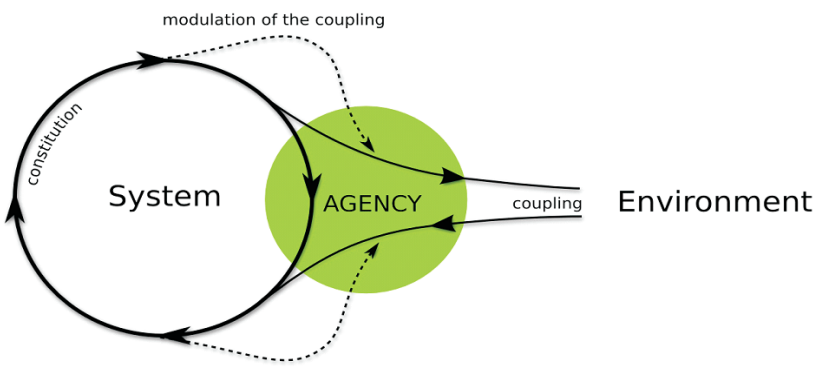

Figure 2.3: Graphical representation of the working definition of agency (copyright © 2009 X.E. Barandiaran [2] under a Creative Commons Attribution Share-Alike license).

In other words, an agent reacts to its environment by predicting future states. Actions of an agent are based on choosing behavior which results in a state which is the most distant from the limits of its viability. Consequently, we can

discriminate between a being which is only able to self-perpetuate and a *living agent*, which also adapts its coupling to the environment. To further distinguish a genuine agent, three conditions have to be met: individuality, interactional asymmetry, and normativity.

The first condition, *interactional asymmetry*, states that the system is the active source of activity in its environment. The coupling between an agent and its environment is an asymmetrical physical happening; an agent is able to modulate some of the parametric conditions and to constrain this coupling in a way that the environment typically does not. In other words, the agent modulates a subset of the conditions and constraints that modulate the coupling (see figure 2.3). Furthermore, an agent is a system that systematically and repeatedly modulates its structural coupling with the environment.

The second condition, *normativity*, states that an agent actively regulates its interaction with the environment and this regulation can produce failure or success according to some norms. These norms are self-imposed, as they are established by the viability conditions of the agent. Specific norms relate to the different ways in which a change in the processes of an agent can lead the agent to lose its organization as a self-maintaning network [2]. In other words, the agent learns by self-reinforcement.

The last condition is *individuality*, which is described as the capability of an agent to define its own identity as an individual and thus distinguishing itself from its surroundings. By this distinction it defines an environment in which it carries out its actions. This condition can be compared with the well-known theory of *embodied cognition* (see [46, 6, 5]). This theory proposes that cognitive processes stem from the real-time interactions between organisms and their environment; the nature of these interactions influences the formation and further specifies the nature of the system. Hence, an agent is able to define itself in an environment, to which it is inherently connected and on which it is dependent.

Metabolism

Metabolism as exhibited by minimal life forms, like one-celled organism, plays a central role in *grounding* the above definitions of biological agency and adaptivity (see [25, 28]). A metabolic system depends on the autocatalytic closure of chemical reactions (where E represents energy, C represents the catalyst, M represents matter, and W represents waste):

$$E + M \xrightarrow{C} C + W \tag{2.5}$$

In a metabolic network (see figure 2.4), energy and matter are lost as heat and waste, requiring continuous acquisition of new resources. Here, the requirement for minimal living organization is a metabolic network of chemical reactions that produces and repairs itself. The existence of the metabolic network relies entirely on the existence of catalyst molecules. In more concrete terms, the energetic flow through the system depends upon the existence of catalysts C which in turn depend upon the flow of energy E through the system [13].

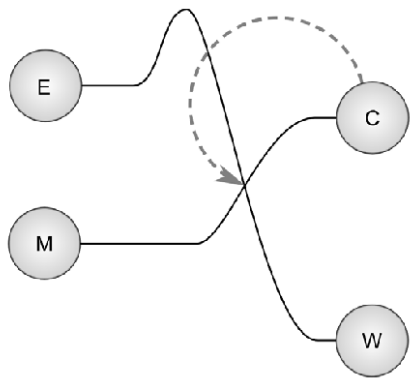

Figure 2.4: Conceptual representation of metabolism, which graphically displays the autocatalytic reaction in equation 2.5 (reproduced from [13]).

The metabolic network adheres to the individuality condition: the system is defined by the interconnectedness of all variables making up the network. Interactional asymmetry is shown by a metabolic system through *chemotaxis*, the movement of the system in reaction to the presence of energy and matter in the environment. Action is defined here as the 'surfing' of environmental effects in the appropriate direction. Here we encounter the normativity condition. Metabolic systems define their optimal environment by the energy and matter present. The metabolism requires the channeling of energy by the catalysts into reactions that produce more of these catalysts. This way, the viability of a system is defined by the number of catalysts present within the system.

Hence, organisms displaying minimal metabolism are systems that must continuously interact with their environment to self-generate and maintain their own precarious organization. Minimal life forms thus satisfy the conditions for agency. This does not imply that all forms of agency need to trace their normative or individuality conditions back to living organization. The essential thing for agency is that, in a manner analogous to that of metabolism, interactive processes can be traced back to a form of organization that displays similar properties [11].

Summarizing the definition of agency, we can say that "[...] for any agentive engagement of a system with its environment its identity must be jeopardized at the proper level and [...] the interaction must involve a process of compensation for deviations from a norm that is generated from within. [...] Actions are guided by the need to compensate the threatening deviation from a norm and environmental processes are integrated into the interaction as relevant for the achievement of such compensation" ([2], p. 378).

When returning to the general definition of agency, the above definitions, especially adaptivity, imply that an agent is able to include a measure of its own viability in the process of generating behavior. The agent regulates its relation to the environment to lead its state away from the limits of viability. Furthermore, because all living agents have finite resources for action, each identified course of action needs to be evaluated in terms of a cost-benefit analysis before it is

initiated.

We will use the conceptual components of this definition of agency as the basis of individual agent behavior. The next section will describe the formalization of these components through the introduction of *affect* as a means of modulating the coupling of an agent with its environment.

2.3 Indivaidual agent behavior

As discussed in the previous section, individual agent behavior can be grounded into the definition of an agent as a self-regulating being. The combination of the two above-mentioned functional demands in the definition of agency, namely access to the agent's viability and resource allocation, have been identified in psychology as the notion of *core affect*, which will be the foundation of our implementation of the individual agent. This section will first describe core affect. After this we will discuss the relation between affect and the motivation of individual behavior.

2.3.1 Core affect

Core affect is "[...] a continuous assessment of ones current state and it affects other psychological processes accordingly" [38]. In his two-dimensional theory of emotion, Russell [38] proposes two primitives that define all other higher-level emotions: *core affect* and *the perception of the affective quality*. The aim of his research is to search for primitive concepts in emotional processes that can exist without intentional objects. Russell defines core affect as the simplest non-reflective feeling evident in moods and emotions. It is proposed as a quantitative emotion theory underlying all basic emotions. As depicted in figure 2.5, it consists of the combination of two dimensions: a hedonic component (pleasure-displeasure), and a motivational component (activation-deactivation).

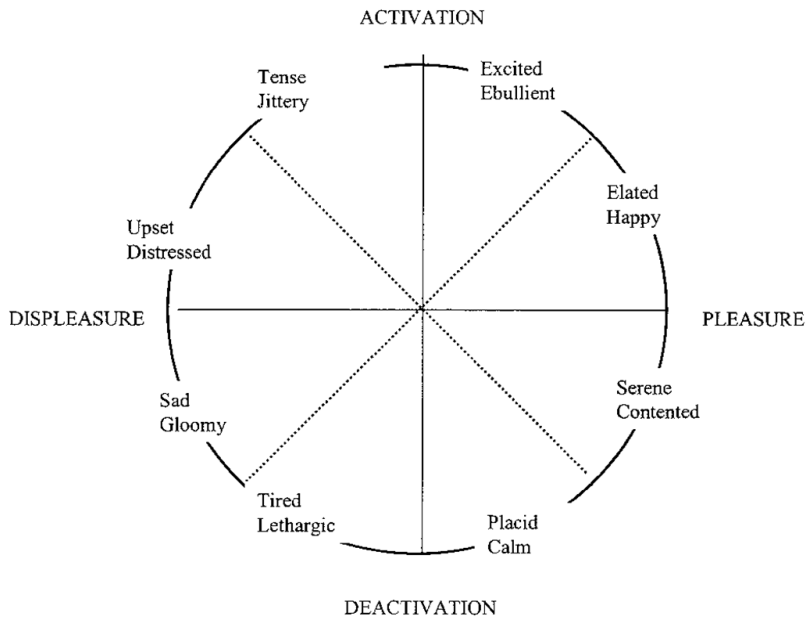

Figure 2.5: The two dimensions of core affect (reproduced from [38]).

These dimensions are experienced without an intentional object (other than the individual self). The affective quality, the capacity to change core affect, is a stimulus-agent combination. The perception of affective quality is a perceptual process that estimates this affective quality. The two primitives define the way in which agents link a change in core affect to its perceived cause. When assuming that people generally seek behavioral options that maximize well-being (or pleasure), Russell proposes that core affect guides cognitive processing and the acquisition of preferences and attitudes. Furthermore, core affect is involved in motivation, reward and reinforcement.

Predictions of future core affect will be used as the main principle for behavior selection in our model. Behavior choices are based on the prediction of future core effect, as follows from the definitions of agency mentioned above. We will call this kind of behavior *affectively motivated*. In the next sections, we will describe current research concerning the relation between emotion and the motivation of behavioral decisions.

2.3.2 Emotion and agency

While we stress that core affect is not to be interpreted as a theory of emotion in the traditional sense, its roots are still firmly based in emotion theory. We will describe current emotion research in order to frame core affect within its foundational theories.

William James [23] (in his version of the James-Lange theory), states that emotions are caused by changes in physiological conditions relating to the autonomic and motor functions. More recently, an emotion was defined as a "[...] collection of responses triggered from parts of the brain to the body, and from

parts of the brain to other parts of the brain" [8], but also as a "[...] vague sensation with uncertain affect pedigree" [39]. Many researchers distinguish between *emotions* (the former quote) and *feelings* (the latter quote).

Core affect stems from the field of *computational* models of emotion [33, 34], in which emotions are defined as quantifiable concepts. This field can roughly be divided into two categories. The first category consists of 'deep models', which describe the situations that initiate the emotions and take into account the constitution of the subjective experience. The second category is made up of 'shallow models', which are about the results of an emotional episode; what follows after an emotion has been experienced? This field can in turn be divided into the sub-fields of *evolutionary models*, *appraisal models*, and *dimensional models*. Core affect falls into the latter sub-field.

We will shortly describe all three kinds of shallow models. First, evolutionary models describe emotions as a result of the selective adaptation to ensure survival. The assumption is that emotion serves as a selection criterion in that it aides, among others, fleeing behavior, sexual attraction, and coalitional aggression (inspired by the domain of evolutionary psychology [3]). It is then assumed that the result of selection is a set of innate, basic emotions.

Second, appraisal models ([39], [33]) are based on the assumption that emotions are the result of an evaluation of interaction between an agent's goals, beliefs, and the current state of the environment. A well-known appraisal model, the *component process model*, predicts that emotions are created and differentiated on the basis of the subjective evaluation of an event on a set of subjective criteria of the agent experiencing the emotion. This model is also based on the distinction between emotions and feelings.

The component process model is based on appraisal theory. This theory predicts that emotions are elicited and differentiated on the basis of the subjective evaluation of an event on a set of *appraisal criteria*. Examples of appraisal criteria (or scales) are novelty, pleasantness, expectedness and coping potential of the stimulus. A feeling state is produced due to the agent's interpretation of the meaning of a stimulus. Scherer and colleagues call this the preconscious appraisal of the intrinsic pleasantness of a stimulus [39].

Finally, dimensional models of emotion can be seen as an abstraction of the aforementioned basic emotions. They are defined as a position in a continuous multi-dimensional space where each dimension stands for a fundamental property common to all emotions. Core affect falls in this range of emotion models.

Frijda [17] makes the division between feelings and emotions on *functional* grounds. He states that feelings refer to the awareness of emotional processes through the environment by which they are caused. An emotion refers to a state of action readiness, which is a motivational state. When an event occurs, we interpret that event: this is our *appraisal*. According to Frijda, the appraisal leads to action readiness, affect and arousal. *These three responses are what motivates behavior*.

Frijda refers to the relationship between behavior and appraisal as *regulation* processes. Previous experiences are evaluated and remembered and influence similar future situations. Appraisal processes are non-conscious (out of awareness) and influence arousal, affect, and action readiness. Hence, appraisal

processes influence future appraisal and action readiness (action tendencies). Habits, or tendencies to behave in certain ways are based on the perceptions of affective states associated with certain situations. As can be seen, this definition of emotions strongly overlaps with the core affect and agency definitions as described in the previous sections. We will use Frijda's definition as a guideline throughout this research.

2.3.3 Emotion in agent-based models

Steunebrink and colleagues [41] provide a logic-based implementation of a BDI agent that exhibits emotional states. They define the incorporation of emotions in a BDI-agent from a technical viewpoint. Agents maintain beliefs about the world, have desires/goals to achieve, employ plans to achieve them, and generate intentions on the basis of these plans. The behavior of BDI agents is described in terms of the evolution of mental states. Emotions moderate the execution and maintenance of the agent's agenda.

Hence, according to Steunebrink and colleagues, emotional agents can be defined as artificial systems that are designed in such a manner that emotions influence decision making. By referring to the research done by Damasio [8], they propose that there is psychological evidence that having emotions may help one to do reasoning and tasks for which rationality seems to be the only factor. Furthermore, citing Picard [34], they propose that emotional states organize ready repertoires of action. In other words, emotions are heuristics.

Emotions can also be used as design tools for an artificial agent architecture. Emotions are used to describe the behavior of intelligent agents; and thus it is useful to reason about the emotional states an agent may be in, and their effects on the agent's actions. For the purposes of our model, the logic Steunebrink and colleagues use is overly focused on the behavior of the system.

For our purposes, the structure of emotions seems to be less important than the function of an emotion as the initiator for behavior. Our model will stray from the above-mentioned designated paths within the field of multi-agent modeling, as it implements decision-making as an implicit function of the perception of the environment. BDI agents reason about the world by directly using evidence from percepts. *Our model is novel in that it models cognition as the regulation of the agent's core affect in relation to the environment.* By this mechanism, belief, desire and intention are self-composed features arising from the agent's internal dynamics.

In the same way as economically rational agents (as defined in section 2.2.1), we will observe and design the affectively motivated agent from the intentional stance. That is, we view the behavior of an agent in terms of mental properties, so that we can reason and make predictions about agent decisions and actions. Hence, by an analysis of the core affect variables over time, we will be able to gain insight into the affective motivation of behavioral decisions. Now that we have discussed the conceptual components of the agent-based model, the next chapter will describe the methodology of the design and testing procedures of the individual affectively motivated agent.

Chapter 3

Model design

In this research, we are interested in answering the following specific research question, as previously stated in chapter 1:

> *How can an affect-based model of human cognition improve the understanding of the motivation of human behavior within a social environment?*

For the model design, this research question boils down to the following, more specific sub-questions:

1. *How does perception of the environment influence the agent state over time?*

   We are interested in how the central driving force of our model, core affect, behaves in individual agents. How is the evaluation of the current emotional state involved in decision-making in social situations?

2. *How does the agent modulate its coupling with the environment?*

   We are interested in the way the agent chooses the (intensity of) actions, and what kind of internal states are the cause of these behavioral decisions. What is the nature of self-regulation in affective agents?

These questions are used to guide the design and verification of the agent-based model. This chapter will describe the global model design and testing procedures in more detail, after which the next chapters will describe the actual model components and experiments.

3.1 Agent design

An individual agent in our framework performs its actions, processes information and perceives the world based on its core affect. The choice of action depends on the coupling between the agent state (actual and predicted core affect values) and the environment state (social space) via perception (social skill). As stated by the definition of agency, an agent acts by modulating this coupling. For our model, we will define the modulation of this coupling as *affect regulation*, as

depicted by figure 3.1. As guiding principle we have that people always aim to optimize viability.

In the proposed framework, affect regulation is conceptualized by the perception of the environment, the internal state of the agent and the action selection. Core affect represents the internal state of the agent. The internal state is influenced by the interaction with other agents in the environment. The agent takes action by choosing a location, which it bases on the predicted core affect. It is important to note that the environment plays an integral part in the determination of the internal dynamics of an agent: environment and agent cannot be separated.

We designed the internal agent architecture in an incremental way. We chose to use this approach mainly because it aides us in the design process and analysis of the final model. By starting of with a simple, basic model, we are able to gain more insight into the workings of the final model.

We first designed a baseline model, which is analogous to the Demographic Prisoner's Dilemma model designed by Epstein [14]. This model exhibits the fundamental elementary structure and dynamics of the world we want to model: the target system. In our case, the target system is an abstract representation of different types of agents in a *social space*. The baseline model also serves as a null hypothesis against which all subsequent extensions can be tested. What do the extensions add, and how do they contribute to more cognitively plausible behavior?

The extensions of the model are represented by adding variables and functions. The eventual goal of this incremental approach is to fully understand our final agent model, which we will call the affective learner agent. For each increment of the model design, we will perform *verification* tests to assess the influence of the changes in the model on the agent state over time.

The first extension (the perceiving agent) of the baseline model represents deliberate *action selection* through the perception of expected future affective states (see chapter 5). Here, all agents are equally proficient in predicting their future affect. The second extension (the learning agent) deals with this cognitively implausible feature by adding a *learning* component (see chapter 6). The implications on the behavior of the agents can subsequently be tested. The final extension (the affective learner agent) of the model adds *motivation* in terms of the investment in action or arousal (see chapter 6). The effects of motivation on the behavior and internal state of the agent is subsequently tested.

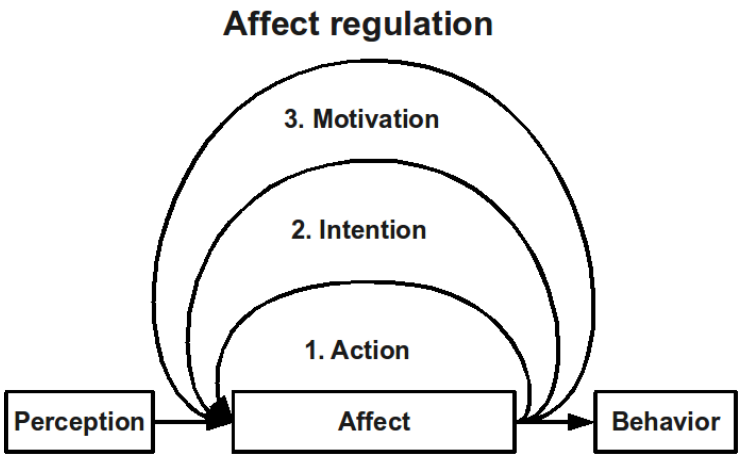

Figure 3.1: Affect regulation lies at the core of agent behavior. It is implemented through the baseline agent and two increasingly complex extensions. The baseline agent implements action, the perceiving agent implements intentional action selection, and the affective learner agent implements motivation of actions through skill learning.

The three levels of affect regulation do also represent the conditions of agency, as described in section 2.2.4. The baseline model only adheres to the condition of interactional asymmetry, since this agent is the source of activity in the environment, but nothing more. The perceiving agent also adheres to the condition of normativity, as it takes deliberate, intentional action according to a norm based on its internal state. Finally, the affective learner agent also adheres to the individuality condition, by adding a motivational component, which modifies the affective influence of interactions through a learning mechanism.

The order of the extensions is also a design choice, as we can now test the influence of learning on agents capable of only action selection and agents capable of both selection and investment in action. The next section will describe the environment in which the different extensions of the individual agent model will be tested.

3.2 Environment

The environment defines the constraints of the world in which the agents live. It defines the rules of the game to which the agents have to adhere, so to say. The environment in which we will test our affective agent framework is analogous to the environment as specified by the Demographic Prisoner's Dilemma (DPD) framework [14].

We have two main reasons to use this framework as the basis of our model. First, the framework is based on the Prisoner's Dilemma payoff matrix, which matches with the subject of our target phenomena, that is, the dynamics among agents in a society of co-operators and defectors. Second, the agents exhibit an internal state, which matches with the definition of an agent as exhibiting an

internal state and the ability to self-regulate through this internal state. Finally, the DPD agents can be used as a baseline to compare the behavior of further extensions of the individual agent model.

The agent society $\{1, .., i, ..., n\}$ is defined by the space on which the agents live, the location of the agents on this space, the interaction between the agents and the current time-step $t \in \mathbb{Z}$. The number of agents n as well as the dimensions can be varied. The space, consisting of discrete points on a *torus*, determines the way in which agents interact. A torus can be projected to a two-dimensional representation by a square lattice with periodic boundary conditions[1] (see figure 3.2).

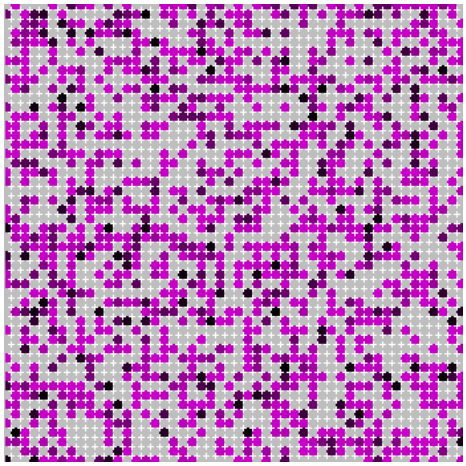

Figure 3.2: The visual representation of a 50x50 torus containing 1225 agents.

Space is especially important when measuring the interaction range I of agents. The location of an agent on the torus determines its interaction range, which is equivalent to the I-Moore neighborhood[2] as illustrated by figure 3.3. Interaction takes place analogously to the DPD model. Each time step, an agent communicates with all the agents within its interaction range. The agent evaluates its internal state and evolves accordingly.

[1] agents which exit the lattice on the left- or top-side, re-enter at the right- or bottom side, respectively.

[2] The eight cells surrounding a central cell on the two-dimensional square lattice.

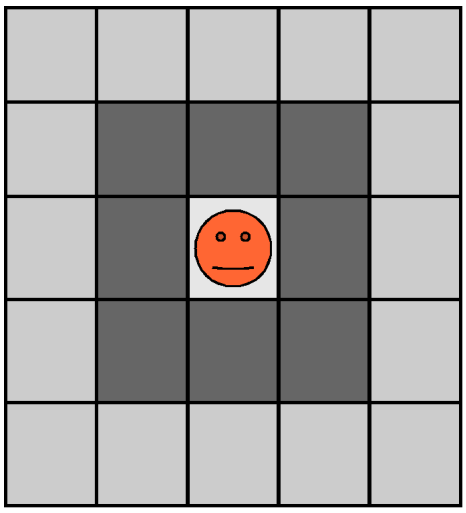

Figure 3.3: The Moore neighborhood of size 1 of an agent is designated by the dark-gray cells. The Moore neighborhood of size 2 is designated by the light-Gray cells.

Time is divided in discrete steps, which represent a certain time period in real life. Agents are updated asynchronously in random order. That is, each time-step, agents are put in a list in random order, and update and make moves serially, one after the other. Asynchronous updating is important, since the simulation has to mimic a real world system without a global clock [22]. A complete run of the model spans multiple generations in real time. Table 3.2 summarizes all world parameters.

symbol	description	domain
t	time-step	$[0,\infty)$
W	torus dimension	$[0,\infty)$
A	proportion of agents	[0,1]
n	number of agents in population	$[1, \infty)$
ϕ	Coid distribution steepness	$[0,\infty]$
I	interaction range	$[0, \frac{W}{2}]$
d_p	payoff matrix dimension	$[2,\infty)$
ψ	death threshold	$(-\infty, 0]$
ρ	reproduction threshold	$[0, 1]$
η	viability transfer	$[0, \rho]$
Y	age limit	[0, 100]

Table 3.1: Description of the world parameters.

Before we go on to the macro-results of the model, the next chapters will describe the implementation of the baseline agent model and its extensions.

Chapter 4

Baseline model

In the previous chapter we have identified the elementary environmental components and the interaction model, which are needed to represent our target world: an abstract society. Now, we need to define the minimal components of an agent in our world needed to be comparable to a living agent. The components of an agent consist of a number of variables and the dynamics between these variables.

The baseline model has the following variables: conscience, payoff and viability (see figure 4.1). The set of functions consists of the action function and the received utility function. The next sections will describe these components.

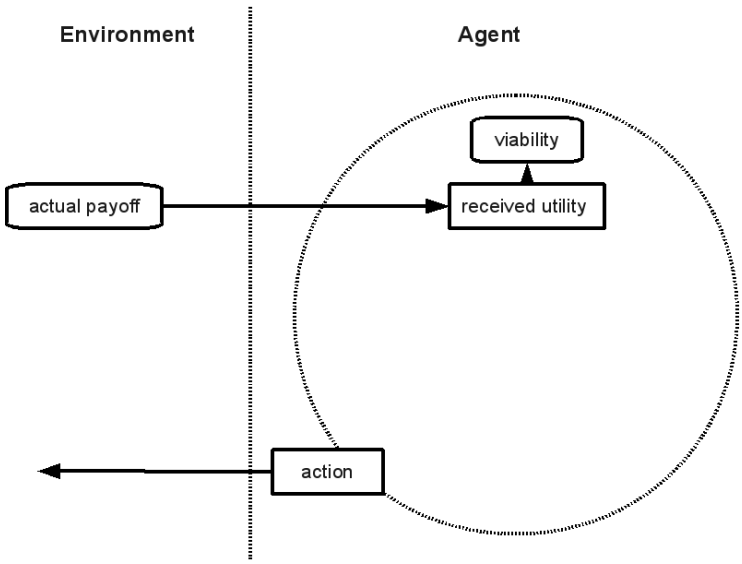

Figure 4.1: The baseline agent.

4.1 Agent components

The components of the baseline agent and the environment as described in the previous chapter serve as the basis for all model increments. This baseline agent formalization will outline the majority of the agent components, while the next three sections will present only a re-implementation or one added component. This section will now discuss the formalization of viability computation function and action function.

4.1.1 Action

In our model, the social situation is represented by the location on the grid. An action is simply represented by the choice of moving to one of the eight squares in the Moore neighborhood, or to remain at the same location. For each time step, the sequence of agents which can make a move is randomly[1] determined. Different agents cannot move to the same location. For the baseline model, the agent takes a random decision out of the nine options. If a position is already occupied, the agent randomly chooses another spot out of the list of remaining locations, and so on.

4.1.2 Conscience

In our model, the conscience $0 < c_i < 1$ designates the genetic foundation for a bias to either cooperative behavior or egoistic behavior. Basically, it determines the *personality* of the agent. It is indigenous to an agent as it is constant through all time-steps. We can initialize the conscience for each agent in two ways. The first is to initialize the agent population is by uniformly distributing the conscience over the agent population. This is done by using the random number generator (as mentioned above) to create the conscience c_i for each agent, where the conscience distribution over the whole agent population is a quasi-normal distribution:

$$c_i \sim \mathcal{N}(\mu = 0.5,\ \sigma = 0.29) \tag{4.1}$$

The second way to initialize the conscience for each agent is by using the distribution as empirically determined by Coid et al. [7] (see chapter 2). From now on, we will call this the Coid distribution. For each agent, conscience is set according to the following conscience initialization function based on the Coid distribution:

$$c_i = -\tanh\left(-\frac{\phi\pi i}{n}\right) \tag{4.2}$$

Where $0 < \phi < \infty$ is a parameter which determines the steepness of the distribution, i is the agent id, and n denotes the number of agents. Hence, the conscience value is a normalized representation of the PCL-R value, as depicted by figure 4.2.

[1] Our implementation uses the Mersenne twister [27] as a pseudo-random number generator.

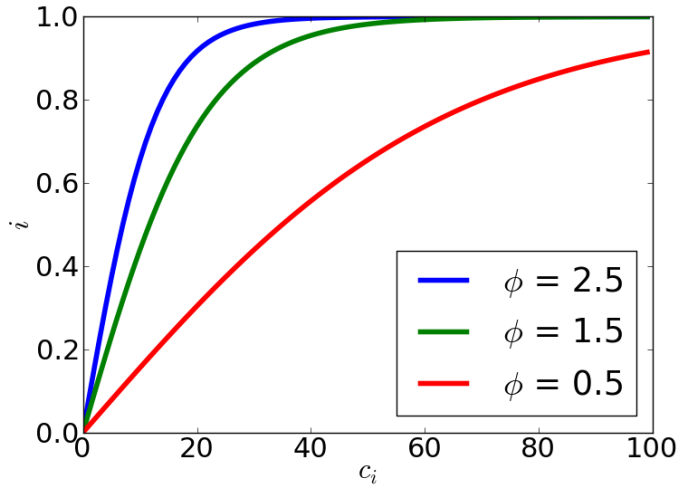

Figure 4.2: The conscience initialization function based on the Coid distribution for a population of $n = 100$ agents and different values of ϕ.

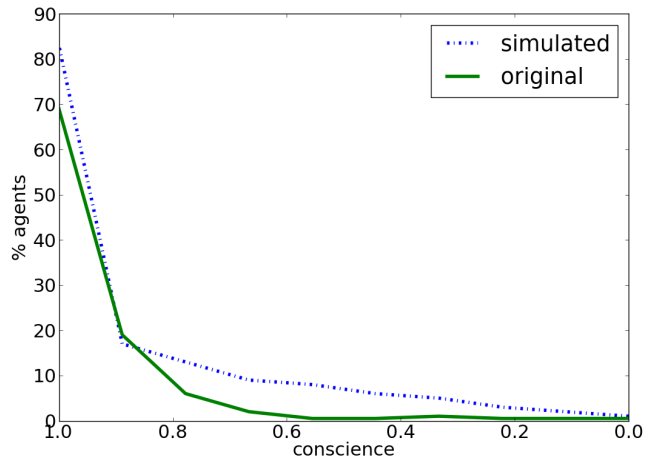

Figure 4.3: Cumulative distribution of the empirically measured and simulated conscience, for $\phi = 1.4$.

By comparing the curves of the simulated and measured cumulative distributions of conscience, we chose a $\phi = 1.4$ as the best fit for the initialization of conscience over the agent population, as shown by figure 4.3.

4.1.3 Received utility

Each time-step, an agent perceives its environment, as defined above. The agent perceives its environment in terms of the average conscience of all agents within the interaction range. The conscience $c_j[t]$ is the conscience of an agent j in the Moore neighborhood at time t.

The payoff matrix fully determines the meaning of the conscience value. On average, low-conscience agents have a negative influence on all agents around them. They do not cooperate, and their aim is to keep the full emotional payoff for themselves. High-conscience agents have a positive influence on the agents around them. By cooperating, they share the affective payoff resulting from an interaction. However, when they encounter a low-conscience agent, they are robbed of their payoff. This entails that high-conscience agents always run the risk that low-conscience agents take advantage of them.

Consequently, a low conscience-agent maximizes the payoff when it can abuse agents in its environment, since it is not influenced by the emotional state of others around it (low empathy). A high conscience-agent maximizes the payoff when it can cooperate (high empathy), and hence lets its environment also receive a positive payoff. The total payoff $p_i[t]$ an agent i receives at one time-step is a linear combination of all the individual payoffs $p_{i,j}$ resulting from playing 2-player games with all n agents in the Moore neighborhood I:

$$p_i[t] = s_I \sum_{j \neq i} p_{i,j} \quad j \in I \tag{4.3}$$

Where $s_I = \frac{1}{8}$ is the size of I. The individual payoff $p_{i,j}$ is computed according to a quasi-continuous payoff matrix based on the comparison between the conscience of agent i and the conscience of agent j. Table 4.1.3 shows the edges of this payoff matrix.

$$p^*_{i,j} = \begin{array}{|c|c|c|} \hline & c_j = high & c_j = low \\ \hline c_i = high & R, R & S, T \\ \hline c_i = low & T, S & P, P \\ \hline \end{array} \tag{4.4}$$

This matrix is homogeneous for all agents. Here, we have to note that the payoff matrix is a square matrix of dimension d_p. The values of the non-edge cells ($0 < i < d_p,\ 0 < j < d_p$) are computed by linear interpolation between the edges. The following conditions hold for the extremes of the payoff matrix: $S < P < R < T$, and $S, P < 0$, and $R, T > 0$ (as discussed in section 2.2.2). In the baseline model, the received utility $r_i[t]$ is proportional to the payoff according to the following equation:

$$r_i[t] = \begin{cases} -E & \text{if } I = \{\} \\ G \cdot p_i[t] & \text{else} \end{cases} \tag{4.5}$$

As can be seen, the payoff is multiplied by the viability change G to achieve the received utility value. This is a fixed parameter for our model.

4.1.4 Viability

The viability $0 < v_i[t] < 1$ of an agent is equal to the resources (energy) an agent exhibits. It is modified by the average influence of the interactions with all the other agents in the current environment. At different timescales it corresponds to self-perceived pleasure, happiness and health. We will say that an agent is able to cope with its environment, when it is able to minimize the negative change in viability and maximize the positive change.

In our model, the viability is initialized randomly within the range $[0, 1]$ and is updated every time-step on the basis of the *received utility* (as described in the previous section). This viability update function computes the updated viability $v_i[t]$ by summing the viability $v_i[t-1]$ and the received utility $r_i[t]$:

$$v_i[t] = v_i[t-1] + r_i[t] \tag{4.6}$$

4.2 Verification

The model as described in the previous section captures the fundamental elementary structure and dynamics of the target system. The most significant features are as follows. First, the social space is represented as a spatial entity. Second, the environment of an agent is determined by the location on the social space. Furthermore, the internal state of the agent is represented by the viability, which can be correlated with the well-being of an agent. Last, each time step, the environment of the agent changes randomly.

symbol	description	domain
y_i	location	([0,d_p-1],[0,d_p-1])
c_i	conscience	[0,1]
v_i	viability	[0,1]
r_i	received utility	[-1,1]
p_i	payoff	[-1,1]

Table 4.1: Description of the baseline agent variables, where d_p is the dimension of the grid.

symbol	description	value
G	viability change rate	0.003

Table 4.2: Agent parameters for the baseline model, constant for all agents.

We will now turn to the analysis of the interplay between the environment and this baseline agent. To compare the dynamics of the baseline model with extended models, we will examine the effect of the environment on the viability of different agent personality types. Are agents able to optimize their well-being when they are behaviorally decoupled from their environment?

Since the baseline agent is not able to take actions in the environment, the configuration of the environment is arbitrary for these experiments. We are only

interested in the effect of the environment on the internal state. We chose to fix the the world dimension to $d_p = 30$ and the agent proportion to $P = 0.1$, as is also done in the DPD model (see [14]). This world configuration is chosen in such a way that the agents have enough space to move to be in a constantly changing environment. The agent parameter and the agent variables are listed in table 4.2 and 4.2), respectively. Furthermore, each experiment uses the same seed to initialize the pseudo-random number generator.

We have the following assumptions about the internal state of the baseline agent, to which it has to adhere. First, the agent's internal variables cannot converge to a static equilibrium. Second, the agent is not able to optimize its viability.

The values R, S, T, P of the payoff matrix and the payoff matrix dimension d_p determine the internal dynamics of the baseline agent. The death threshold ψ and the reproduction threshold ρ determine the time scale of the model, as it specifies the rate by which agents update the internal state, here solely represented by the viability. The edges of the payoff matrix are given by the DPD matrix:

$$p^*_{i,j} = \begin{array}{|c|c|c|} \hline & c_j = high & c_j = low \\ \hline c_i = high & 5, 5 & -6, 6 \\ \hline c_i = low & 6, -6 & -5, -5 \\ \hline \end{array} \tag{4.7}$$

Figure 4.4 shows the received utility for different agents over time plotted against the number of cooperators and the number of sociopath. Here, the agent society is initialized according to a uniform distribution. All values x in the time series are normalized to z-values according to:

$$z = \frac{x - \mu}{\sigma} \tag{4.8}$$

Note that range of z-values differs between the panels shown in figure 4.4. This poses no problem, as we are only interested in the *relative* effect of different kinds of environments on the viability dynamics of a single agent.

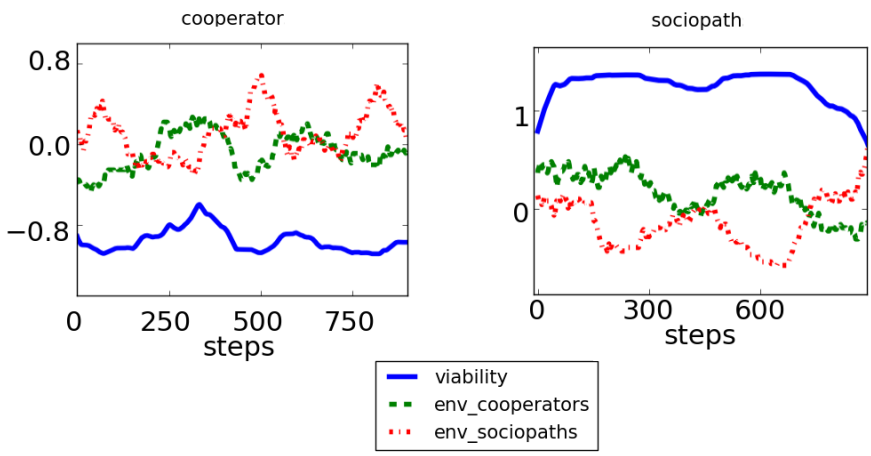

Figure 4.4: Influence of the environment on the received utility for two baseline agents for 1000 steps. Both panels compare the normalized number of cooperators (green), the number of sociopaths (red) and the viability (blue), for a cooperative agent (left panel) and a sociopathic agent (right panel).

Figure 4.4 (left panel) shows the viability of a cooperative agent. We can observe peaks in viability ($t = 350$ and $t = 550$) when the z-value of cooperators in the environment is higher than the z-value of sociopaths in the environment. The sociopathic agent (right panel) shows analogous behavior: the viability is constantly on a high level. Until $t = 750$, there are more cooperators than sociopaths in the environment. After this time-step, the environment of this agent contains a majority of sociopaths. Following this, the viability drops.

These results are intuitively explainable according to the payoff matrix in table 4.7. As said, the payoff matrix states that, for individual encounters, both cooperators and sociopaths receive positive payoff when interacting with other cooperators, while they receive negative payoff when interacting with sociopaths. Following this, the main trend for the cooperative agent (see figure 4.4, left panel) is the negative influence on the viability change of sociopaths in the environment.

Hence, the viability dynamics of an individual agent rely on the ratio of cooperators and sociopaths in the environment. This is true for both the sociopathic and cooperative agents. The viability change parameter G also influences the viability dynamics. A higher viability change value will cause more erratic viability dynamics, since the viability change is amplified (see equation 4.5.

Concluding, the 'well-being' of the baseline agent is fully dependent kind of environment it lives in. This is caused by the fact that a baseline agent is not able to modulate its behavior. When looking at the definition of biological agency (see section 2.2.4, the baseline agent only adheres to the condition of interactional asymmetry. That is, the agent is the source of activity in the environment. This activity is, however, not based on any regulation mechanism, i.e. action is not intentional or deliberate.

The next section will present the first extension of the baseline model. This

extension will add deliberate behavior to the baseline model. In this way, the agent can adhere to the second condition of biological agency: normativity. The agent will choose its actions according to some self-imposed norm. By this mechanism, an agent will be able to choose the environment they live in, and will therefore be able to implicitly self-regulate its viability.

Chapter 5

First extension: the perceiving agent

As the first extension of the model, we will implement action selection according to the perception of expected utility. Hence, an agent chooses to which location it will move according to the expected affective payoff associated with this location. In this way, the agent is fully coupled to its environment. Furthermore, the agent can adhere to the second condition of biological agency (see section 2.2.4). Figure 5.1 shows the added components and the dynamics between variables.

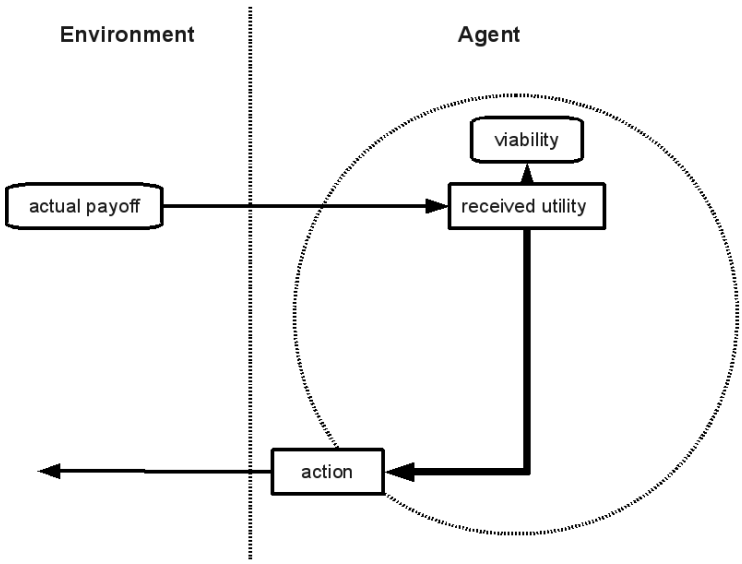

Figure 5.1: The relation among agent components in the first extension of the baseline model: the perceiving agent. This model adds deliberate action through perception of future affective states. Novel components are bold-faced.

5.1 Agent components

The components of the perceiving agent will extend the baseline model with deliberate action, as opposed to the baseline model, in which the future location was chosen randomly. This section will now describe action selection through the received utility value.

5.1.1 Perceiving the environment

For each interaction, the expected utility function predicts the future utility from this interaction according to the observed conscience of the other agent it is interacting with. In our model, agents aim to select an action which results in the optimal utility (increase in happiness). An agent can only observe its direct Moore neighborhood. When the agent moves, in essence, it choose to go to a new environment, as depicted by figure 5.2. The 1-Moore neighborhood is an assumption for our model; we chose to use this approach because of its ease of implementation and low computational complexity.

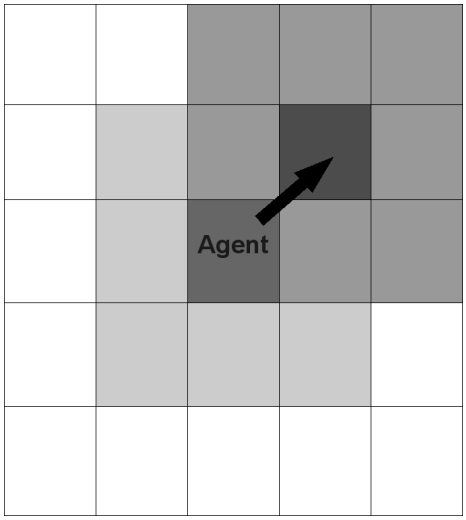

Figure 5.2: The 2-Moore neighborhood of an agent.

5.1.2 Action

People always try to be in the situation in which they expect they will be most happy [38], where viability is a measure for the level of happiness. This is also the case for the first extension of the baseline agent. It selects an action according to the received utility value:

$$\begin{cases} \text{move} & \text{if } r_i <= 0, \\ \text{stay} & \text{if } r_i > 0 \end{cases} \tag{5.1}$$

Hence, the perceiving agent does not want to be in a situation in which it becomes less 'happy'. Hence, unhappy agents move, and happy agents stay

in the same place. Agents are again randomly selected to take turns. If the location is already occupied, the agent chooses the action corresponding to the second highest expected utility value, and so on.

5.2 Verification

When compared to the baseline agent, the perceiving agent is now coupled to the environment, since it is able to take actions according to the expected future internal state. This section describes the behavior stemming from the addition of deliberate action as an extension to the baseline agent.

symbol	description	domain
y_i	location	([0,d_p-1],[0,d_p-1])
c_i	conscience	[0,1]
v_i	viability	[0,1]
r_i	received utility	[-1,1]
p_i	payoff	[-1,1]

Table 5.1: Description of the perceiving agent variables, where d_p is the dimension of the grid.

symbol	description	value
G	viability change rate	0.003

Table 5.2: Agent parameters for the perceiving agent model, constant for all agents.

As mentioned in the previous section, the baseline agent was essentially decoupled from the environment, consequently the configuration of the environment was separable from this agent. The perceiving agent however, is dependent on its environment, since its decisions affect its own internal state *via* the environment. In this section, we are interested in the behavior of the perceiving agent within different environments. Tables 5.1 and 5.2 list the values of the components of the perceiving agent.

We have the following assumptions about the behavior of the perceiving agent. First, most agents are able to optimize their viability by forming clusters (a spatial equilibrium). Second, the agent's internal variables *can* reach a static equilibrium.

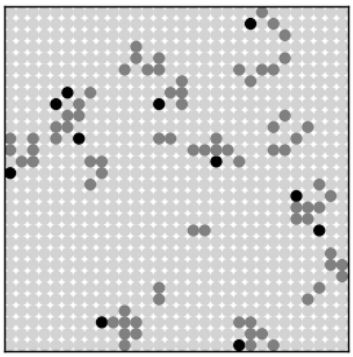

Figure 5.3: Spatial configuration of the perceiving agent population at $t = 747$, after the spatial equilibrium has set in. Light-gray agents are cooperators, dark-gray agents are sociopaths.

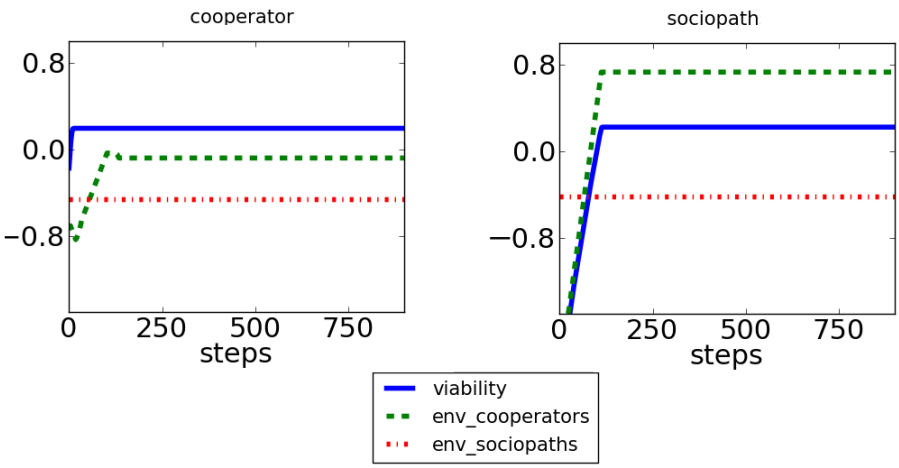

Figure 5.4: Influence of the environment on the received utility for two perceiving agents for 1000 steps. Both panels compare the normalized number of cooperators (green), the number of sociopaths (red) and the viability (blue), for a cooperative agent (left panel) and a sociopathic agent (right panel).

Figure 5.3 shows the spatial configuration of a population of perceiving agents. As can be seen, all high-conscience agents (light-colored) tend to cluster together, while the low-conscience agents (dark-colored) are on the outside of these clusters. The high-conscience agents are able to stay inside clusters equilibrium, while the low-conscience agents have to travel from cluster to cluster. Clusters are sustained throughout the run, and eventually all agents stay on the same location; they are all able to optimize their viability.

Figure 5.4 shows the behavior of the viability over time for two individual agents. The panel on the left shows a cooperating agent in a cluster. Here, we observe that the viability of this agent quickly converges to an equilibrium, which

is caused by the stable environment it exists in. The viability dynamics of the individual sociopathic shows the same prototypical features. The environment quickly converges to a static equilibrium, and the viability stays on the same level throughout the run.

Overall, we see that more stable environments arise in a society of perceiving agents. As was described in this section, the only addition to the baseline model is deliberate action based on the internal state. We have to note that the perceiving agent always makes perfect decisions, based on the limited information it receives about its direct Moore neighborhood. Each agent chooses a neighborhood in which the viability change is positive (see equation 5.1), and stays there throughout the run.

Concluding, all perceiving agents are 'happy' throughout their life. As will be obvious, perceiving agents which optimize their viability and do not take action throughout most of their lives cannot serve as comparison to the behavior of biological or human individuals. Moreover, the perceiving agent does not yet fully comply to the definition of biological agency (see section 2.2.4. To let our agent adhere to the condition of individuality, as the system has to define itself by the interconnectedness of all variables within its internal network. For this, we need a feedback mechanism in our model, which connects all internal variables. This feedback can be provided in the form of a motivational component based on a learning mechanism. The next section will add this component as the second and final extension of the baseline model.

Chapter 6

Second extension: The affective learner agent

By the previous extension, the perceiving agent, all agents are able to make behavioral decisions. However, they are *all* excellent in predicting their future viability. An agent capable of selecting the optimal location, is also able to optimize its viability. The agent is thus able to decide which course of action to undertake. However, it cannot decide on the *intensity* of its action. How is the intensity of the action connected to the quality of the perception of future emotional payoff?

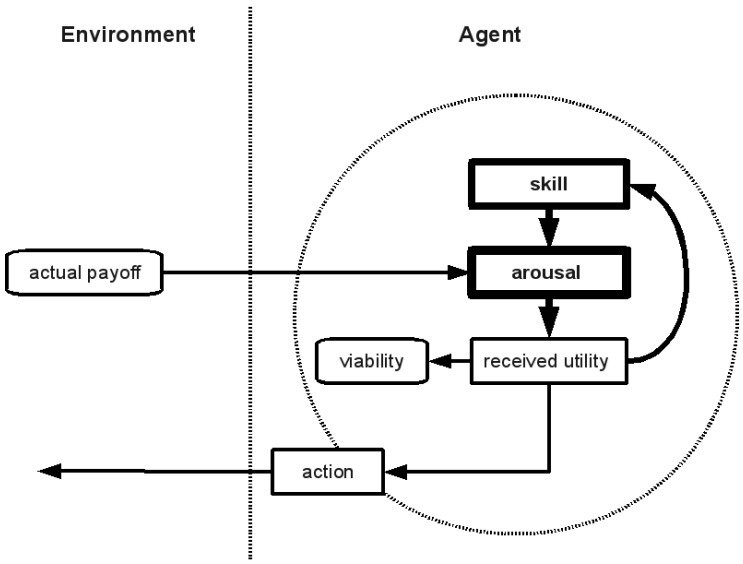

Figure 6.1: The relation among agent components in the second extension of the baseline model: the affective learner agent. This model adds deliberate action through perception of future affective states. Novel components are bold-faced.

The above-mentioned question leads us to the definition of *arousal* within the core affect theory [38], as described in section 2.3.1. Heightened arousal contributes to the intensity by which any action plan is executed. In the second extension of the model, we formalize investment in action according to the prediction of future viability change.

We will add two variables to the individual agent model: the *social skill* variable and the associated *arousal* variable. Figure 6.1 shows the relation between arousal and the other attributes of the agent. Below, we will formalize the function of learning through arousal choice.

6.1 Agent components

The components of the learning agent will extend the perceiving agent by introducing quality of perception proportional to the novel skill variable. The arousal is formalized as a modifier for the amount of payoff. This section will describe the formalization of this extension by the function which computes the utility according to the expected payoff (based on the skill variable) and arousal values for each interaction.

This section will describe a number of new components. The first is the skill update function, the second is the perceived payoff function, and the third is action selection through the perceived utility value.

6.1.1 Social skill

The social skill level $0 < z_i[t] < 1$ is a variable that represents the *social intelligence* of an agent. As reviewed in chapter 2, we interpret the social skill as 'what to do when'. In our model, a higher social skill level entails that an agent is better in predicting the expected utility vector $\vec{x}_i$. Consequently, the social skill level determines the degree to which the agent makes optimal behavioral decisions.

In the perceiving agent model, all agents make the relatively optimal decision. We have to note that they cannot make absolutely optimal decisions, since they do not take the future location of the neighboring agents into account; they are not capable of predicting future decisions of other agents. For our purpose, the social skill discriminates agents by the quality of their actions. Along with the fixed conscience value, it determines the changing strategy quality (degree of adaptation to the environment) of an agent throughout its life.

6.1.2 Skill update

Organisms learn by making mistakes, by evaluating the *quality* of their (inter)actions. In our model, the social skill value is modified based on the perception of the quality of the previous action. We use the Temporal Difference (TD) learning algorithm [45] as basis for our skill update function. The TD algorithm is a reinforcement learning algorithm which uses the discrepancy between the actual reinforcement and the last prediction of the reinforcement as the basis for learning. The reinforcement in our model is the affective impact of interactions (or the received utility) during one time-step.

The TD algorithm is ideally suited for our purposes, since the agents evaluate the quality of an action based on the difference between the *actual* affective impact of an action and the *expected* affective impact. Hence, the agent learns by self-reinforcement. We can formalize the algorithm, by computing the difference between the received utility and the expected utility, in the following way:

$$z_i[t] = z_i[t-1] + L(r[t-1] - x[t]) \tag{6.1}$$

Here the parameter L controls the skill learning rate. The actual skill value is updated according to the normalization function (see also equation 6.3):

$$z_i[t] = g(z_i[t]) \tag{6.2}$$

The normalization function is a logistic function and has the following general form (table 6.1 shows the fixed parameters of this function):

$$g(k) = s_p(\tanh(s_p\pi(k + s_n)) + s_q) \tag{6.3}$$

s_m	range position	0.5
s_q	range shape	1
s_m	domain position	2
s_n	domain shape	-1

Table 6.1: Parameters which determine the shape of the logistic normalization function.

Hence, agents learn by making mistakes. Learning does not depend on the sign of the received utility, only on the quality of perception of the expected utility value. The skill increases by the skill change parameter value when the affective impact of interaction was overestimated or underestimated. The skill decreases by the skill decay parameter value when the the affective impact was as expected.

6.1.3 Expected utility

As mentioned, the social skill value influences the quality of perception. We introduce a skill threshold δ, which determines the quality by which the agent perceives the expected utility of an interaction.

$$x_i[t] = \begin{cases} -p_i[t+1] & \text{if } z_i > \delta, \\ p_i[t+1] & \text{if } z_i <= \delta \end{cases} \tag{6.4}$$

The skill threshold δ determines when an agent makes an underestimation or an overestimation of the actual payoff. Hence, the probability of 'noisy' observations is reciprocally proportional to the skill level.

6.1.4 Arousal

The arousal $0 < a_i[t] < 1$ is defined as the resource invested in each interaction. Because all living agents have finite resources for action, resource management is essential and each identified course of action needs to be evaluated in terms of a cost-benefit analysis before it is initiated. Following the definition of core affect [38], arousal can be seen as the amount of resources available for immediate action.

In our model, the arousal is a measure of the expected change of the situation of the agent. In this way, it is directly linked to the expected utility; the arousal value is based on the expected utility value associated with an interaction.

The arousal is a normalized representation of the expected utility value associated with the selected location:

$$a_i[t] = g(x_i[t]) \tag{6.5}$$

The equations in this and above subsections fully reflect the way in which the agent selects its behavior. The agent first predicts which action will provide the optimal change in viability, by 'simulating' the action. When it has chosen an action, it makes the action 'hot', by investing a certain amount of affective resources, reflected by its arousal.

6.1.5 Received utility

The more an agent expects to improve its situation, the higher its motivation to take action. Hence, an agent will invest viability in an action, proportionally to the amount of payoff it expects. The received utility is thus enforced by the affective investment in the action (arousal). If the investment in action is large, the agent takes a higher risk. The arousal is a modifier of the eventual received utility, by multiplying it with the payoff value. This leads us to the following extended equation for the received utility:

$$r_i[t] = G \cdot p_i[t] \cdot a_i[t-1] \tag{6.6}$$

From this equation, it follows that if an agent expects an improvement (high arousal), and the actual situation is deteriorated (negative payoff), the viability decrease is 'amplified'. If an agent expects a deterioration (low arousal), and the actual situation is deteriorated (negative payoff), the viability increase is 'dampened'. This way, through arousal, the agent can regulate by which degree the viability change is influenced by the affective impact of interactions. The agent invests viability in action through arousal. A high-skilled agent will typically optimize the payoff on average by investing the right amount of arousal, while a low-skilled agent will show the opposite behavior, by detrimental combinations of arousal and payoff.

6.2 Verification

This section formalized the second and final extension of the baseline model, which fully implements the theory of core affect. We can now describe the algorithm of the full model: the affective learner agent. First, we describe the initialization algorithm of the agent:

```
for agent number do
  create agent
  rec utility = 0
  viability = viability transfer
  skill = 0.5
  age = random
  conscience = coid
end for
```

After the initialization, agents are called asynchronously. For each agent, the functions are called serially. The functions consist of an *action* phase, a *perception* phase, and an *evaluation* phase.

```
for each step do
  set random move order
  for all agents in move order do
    choose action
    move
    for all other agents in environment do
      update arousal
      interact
    end for
    update rec utility
    update viability
    update skill
  end for
end for
```

The following tables list all the variables present in the model. Table 6.2 lists the parameters for the agent functions.

symbol	description	domain
y_i	location	([0,d_p-1],[0,d_p-1])
c_i	conscience	[0,1]
$z_i[t]$	skill	[0,1]
$v_i[t]$	viability	[0,1]
$r_i[t]$	received utility	[-1,1]
$\vec{x}_i[t]$	expected utility vector	[-1,1]
$p_i[t]$	payoff	[-1,1]

Table 6.2: Description of the agent variables, where d_p is the dimension of the torus.

symbol	description	value
G	viability change rate	0.003
D	skill decay	0.003
L	skill learning rate	0.04

Table 6.3: Agent parameters, constant for all agents.

The assumptions for the internal state of the learning agent are as follows. First, agent viability and social skill level are correlated; viability changes proportionally to social skill. Second, agents have to keep making mistakes to optimize their viability; they do not learn from habits. Third, the agent is able to decrease the negative influence and increase the positive influence according to its investment in action. Finally, we expect that the skill value of the agent influences the viability implicitly by the quality of the arousal value.

In this section, we verify the behavior of the skill update function and its implicit influence on viability via arousal choice. Note here that the skill learning function and the arousal choice is equal for both agent types. Hence, the internal influence of the received utility, skill and viability on each other is analogous among all agent types. Figure 6.2 shows the interplay between the received utility and the skill value over time for a (cooperative) individual affective learner.

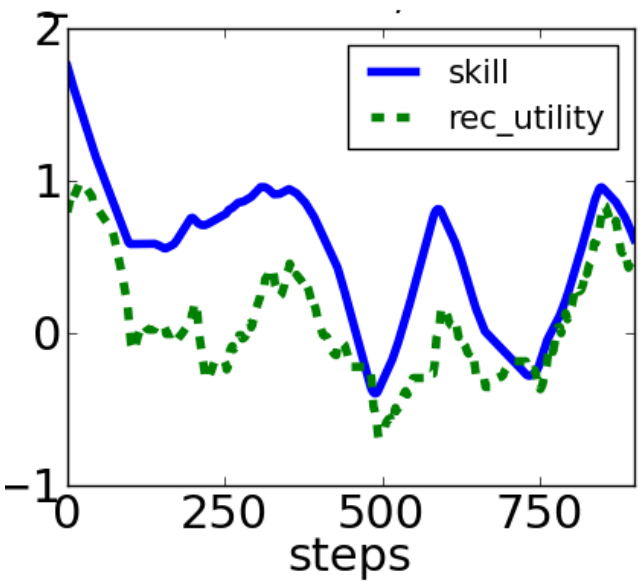

Figure 6.2: Behavior of skill learning by the influence of received utility values for a cooperative agent through 1000 steps.

The main observation of figure 6.2 is the influence of the received utility value on the skill change. When the received utility value drops under a certain value (which is not identified, since the values are normalized), the agent will learn according to the skill learning rate L. This explains the sharp drops in the skill value on $t = 500$ and $t = 750$. These dynamics do also work the other way around. While the received utility has a *explicit* influence on the skill change through the skill update function, the skill value has an *implicit* influence on the received utility through the arousal choice (see equation 6.6). As the skill decreases, the received utility will also decrease. Figure 6.2 shows this effect between $t = 250$ and $t = 500$, where the agent first has a relatively high skill, which results in a higher received utility. From this, the skill decreases by the skill decay rate D. Almost instantly, the received utility also decreases, and both variables reach a minimum at $t = 500$. Note here that the negative change in received utility does not occur directly, since it only modulates the influence of the environment. Hence, the ratio of agent types in the environment, together with the modulation of affective impact, determines the individual agent dynamics.

Figure 6.3 shows the time series of the skill and viability values. Until $t = 750$, we can observe that the skill value is almost fully out of phase with the viability value. This phenomenon can be explained by distinguishing the features of the self-reinforcement of the viability value through the feedback loop between the skill and viability values via the arousal choice. The loop starts when the received utility becomes negative. Two things occur: first, the agent starts to move, and second, the agent starts to learn by the skill learning rate L (see equation 6.1). This entails an increase in the skill value. After some time, the now high-skilled agent stops moving as its viability increases through the

correct modulation of the positive influence of the environment (see figure 6.6). The skill now decreases by the skill decay rate D, since the viability increases, and the received utility is positive. After some time, the agent will perform detrimental modulation of the affective influence of the environment, which entails a drop in the viability values. The received utility becomes negative, and the loop will restart.

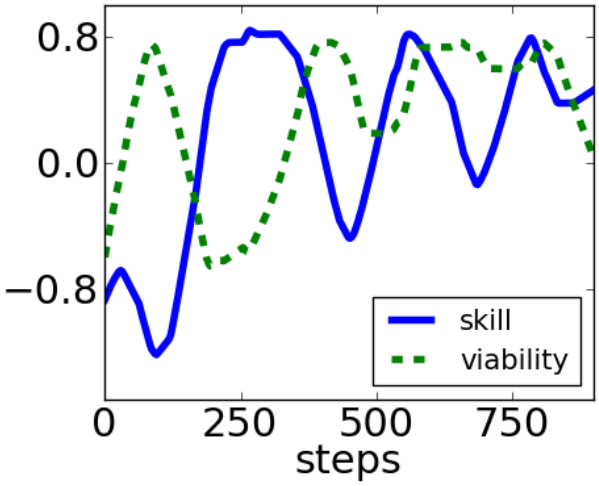

Figure 6.3: The interplay between skill and viability values for a cooperative affective learner agent over 1000 steps.

Concluding, the viability dynamics of an individual affective learner agent are inherently of an oscillating nature. A high viability value is caused by a situation in which the agent does not learn, which eventually causes the viability to decrease by the detrimental motivational choices a low-skilled agent makes. It is important to note that the skill influences the viability implicitly via the modulation of the affective impact of the environment. A low-skilled agent which is not confronted with sociopaths does not get unhappy, but does not learn either.

We have now fully formalized and analyzed the behavior of the individual affective learner agent. This agent adheres to the three conditions of biological agency, as described in section 2.2.4, while also implementing an interaction model for the application in social simulations. The next section will analyze the macro-social dynamics among the agent implementations described in this section and the last two sections.

Chapter 7

Validation experiments

This chapter will show the results of macro-social phenomena arising from multiple runs of the baseline model and its extensions. We will investigate the influence of interaction and spatial configuration among agents on the health (or well-being) of the whole population. We describe two experiments. In the first experiment we will investigate the behavior of a population of agents without evolution. The second experiment is an attempt to introduce a truly precarious environment, by adding the possibility of agent death and reproduction. We compare the results of the addition of this evolution condition with the results of the first two experiments of the Demographic Prisoner's Dilemma (DPD) by Epstein [14].

This chapter will focus on the following general experimental research question.

- *What are the affective factors that drive the dynamics of cooperation in a society of agents?*

We have a number of general goals for this chapter. First, we propose that the simulation of a society of interacting affective learner agents is sufficient to generate a histeroidal cycle.

Second, we propose that a histeroidal cycle as generated by a society of affective learners can be seen as the low-level dynamics underlying the dynamics of cooperation. We hypothesize that the DPD paradigm presents too much a simplification of reality to explain the dynamics of cooperation as it cancels out the histeroidal cycle.

Third, we propose that the addition of punishment to the model re-instates the histeroidal cycle. We propose that this final model can provide an affect-based explanation for the dynamics of cooperation and that this is a more cognitively plausible explanation than the DPD model can provide.

7.1 Experiment 1: baseline model and extensions without evolution

This section will describe the first experiment, in which evolution is left out. We chose to leave out evolution so that we can analyze the macro-dynamics stemming directly from the interactions between agents without having to deal with the added complexity of reproduction and dying. Tables 7.1 lists the world parameters for this experiment. In a world without evolution, the reproduction and death thresholds are used to normalize the values of viability (v_i), skill (z_i) and arousal (a_i). Normalization is applied by using a logistic function which keeps all values between the thresholds.

symbol	description	value
d_p	dimension	30
A	proportion of agents	$\frac{1}{9}$
n	number of agents in population	100
ϕ	Coid distribution steepness	1.4
I	interaction range	1
d_p	payoff matrix dimension	2
ψ	death threshold	0
ρ	reproduction threshold	1
η	viability transfer	0.6
Y	age limit	100000

Table 7.1: World parameters. Changes from these values will be mentioned in the text.

symbol	description	value
G	viability change rate	0.003
L	skill learning rate	0.04
C	skill decay rate	0.003

Table 7.2: Agent parameters, equal for all agents. Changes from these values will be mentioned in the text.

The agent parameters (see table 7.2) list the parameters for the individual agent functions as described in chapters 4, 5, and 6. Note here that for the baseline and perceiving model, only the G parameter is used.

7.1.1 Baseline agent

The baseline model represents an agent which can only adhere to one condition of biological agency: interactional asymmetry, since the agent is the source of activity in the environment. We could even dispute this ruling, since the the baseline agent takes random actions. The baseline model will be used in order to compare its macro-behavior with its extensions.

For the baseline agent, we perform the experiments using two different initialization functions for the agent conscience (see section 4.1.2). Firstly, we test a world in which the agent conscience is uniformly distributed, hence where the conscience is randomly initialized for each agent. Secondly, we test a world in which we use the sociologically plausible Coid distribution to initialize the conscience of each agent. We have the following hypotheses for the macro-dynamics of the baseline agent model. First, mean viability values do not show periodic behavior, since behavior is decoupled from perception. Second, we do not expect a relation between the viability of sociopaths and cooperators over time.

Uniform distribution

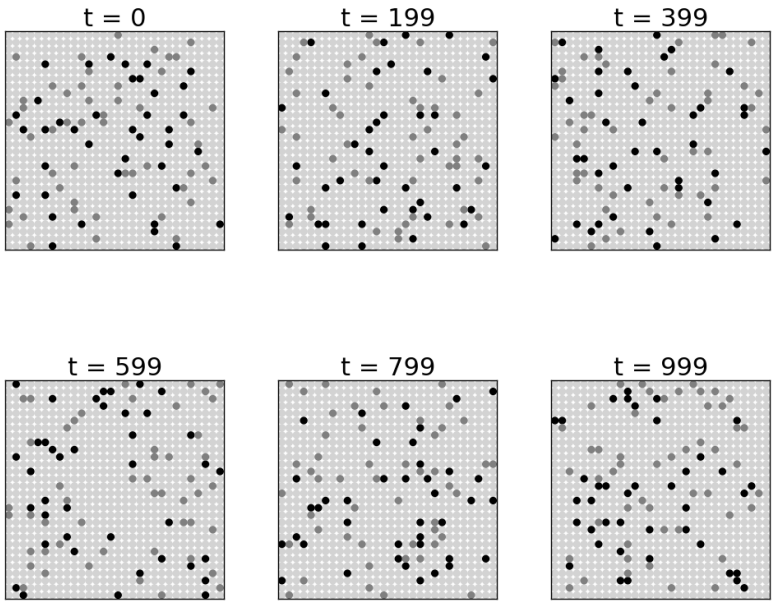

Figure 7.1: Snapshots of the spatial configuration at equally spaced time-steps t for a run of the baseline agent population over 1000 steps, where the agent society is initialized according to the uniform distribution.

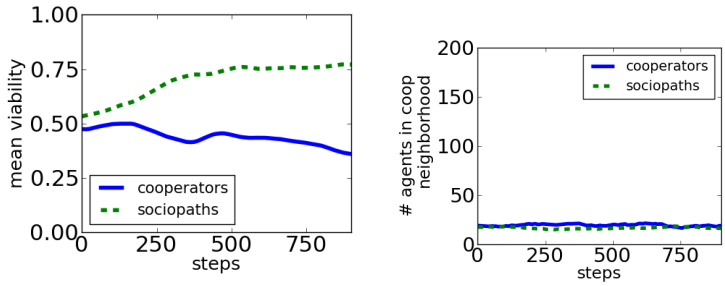

(a) (b)

Figure 7.2: A comparison between the viability dynamics of a society of baseline agents and the mean neighborhood distribution for the cooperative baseline agents over 1000 steps. The left panel compares the mean viability for cooperative and sociopathic baseline agents. The right panel shows the degree of clustering of cooperative agents by showing the number of agents of each agent type in the neighborhood of *cooperative* agents. The conscience for the agent society is initialized according to the uniform distribution.

Coid distribution

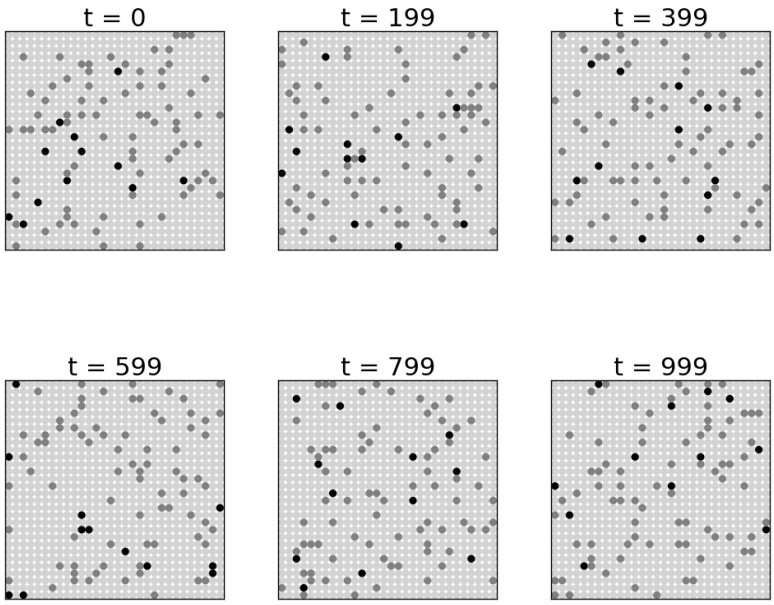

Figure 7.3: Snapshots of the spatial configuration at equally spaced time-steps t for a run of the baseline agent population over 1000 steps, where the agent society is initialized according to the simulated Coid distribution, where $\phi = 1.4$.

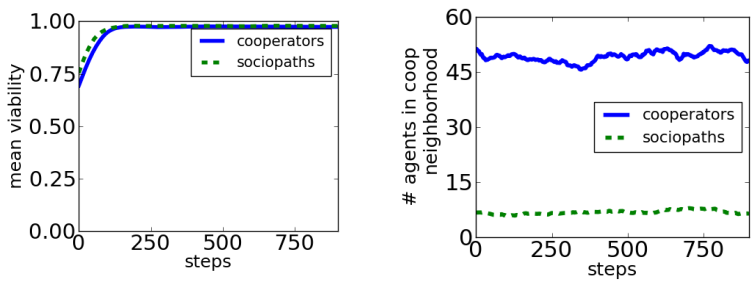

(a) (b)

Figure 7.4: A comparison between the viability dynamics of a society of baseline agents and the mean neighborhood distribution for the cooperative baseline agents over 1000 steps. The left panel compares the mean viability for cooperative and sociopathic baseline agents. The right panel shows the degree of clustering of cooperative agents by showing the number of agents of each agent type in the neighborhood of *cooperative* agents. The conscience for the agent society is initialized according to the uniform distribution. The conscience distribution is based on a simulated Coid distribution, where $\phi = 1.4$.

Discussion

From figures 7.1 and 7.3, we see that, throughout the run, agents are randomly dispersed throughout the torus. No stable clusters exist throughout the run, since agents do not show deliberate action. and we cannot distinguish periodic movements.

Figure 7.2(a) shows the viability dynamics for the two populations when they both have an equal number of agents in the population. In this world, where agent types are uniformly initialized, sociopaths dominate, like is shown in the traditional Prisoner's Dilemma (as discussed in section 2.2.1). This uniform distribution of agent types exactly adheres to the condition set by the payoff matrix for the baseline agent, as depicted in table 4.7, where the sociopath maximum and minimum payoffs are higher than the cooperator maximum and minimum payoffs, respectively.

In the real world, society does not consist of an equal number of sociopaths and cooperators. This is why we test the effect on the mean viability dynamics for a skewed distribution of agent types. Figure 7.4(a) shows the viability dynamics for the two populations when they are initialized according to the more sociologically plausible Coid distribution (see section 2.1.1).

The Coid distribution shows a convergence of both populations to an optimal mean viability. We attribute this to the fact that both cooperators and sociopaths do encounter mostly other cooperators. This causes both of the agent types to optimize their viability. This fact is illustrated when comparing the total number of agents surrounding cooperators between the conditions, as depicted by figures 7.2(b) and 7.4(b).

We can now conclude that in a non-evolving society of baseline agents, in which a small minority of sociopaths is present, the society well-being is always stable at the highest level. The next experiments will show whether the extensions of the baseline agent do show different mean viability dynamics for both agent types.

7.1.2 Perceiving agent

The perceiving model represents an agent which adheres to two conditions of biological agency (see section 2.2.4). Agents are both the source of activity in the environment (interactional asymmetry), while they also evaluate the affective result of their actions. Hence, perceiving agents also adhere to the condition of *normativity*.

We have the following hypotheses for the macro-dynamics of the perceiving agent model. First, we expect the cooperators to form clusters, since they will stick to environments in which the viability change is positive. Second, we expect the sociopaths to engage the clusters of cooperators. Figures 7.6(a) and 7.5 show the dynamics of mean viability and the spatial distribution of the population of perceiving agents. We used the Coid distribution (see 2.1.1) to initialize the agent population, since this distribution best reflects the real-world situation.

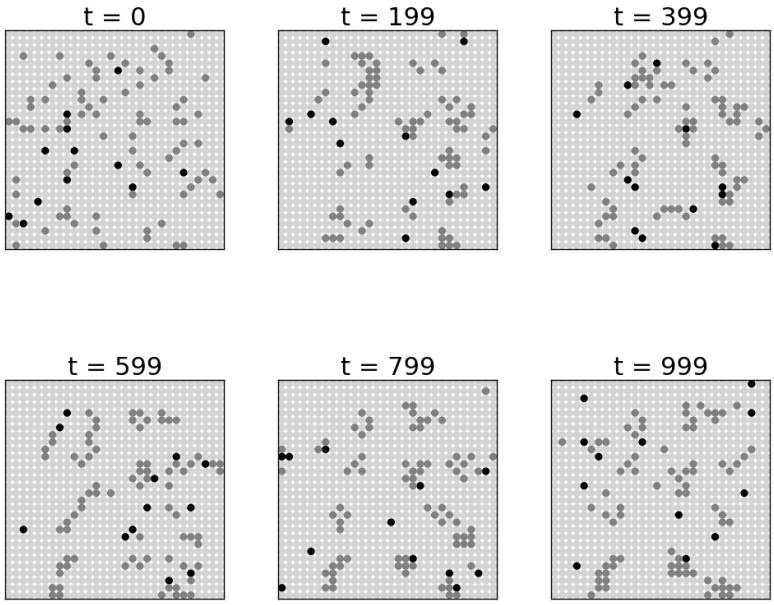

Figure 7.5: Snapshots of the spatial configuration at equally spaced time-steps t for a run of the perceiving agent population over 1000 steps, where the agent society is initialized according to the simulated Coid distribution, where $\phi = 1.4$.

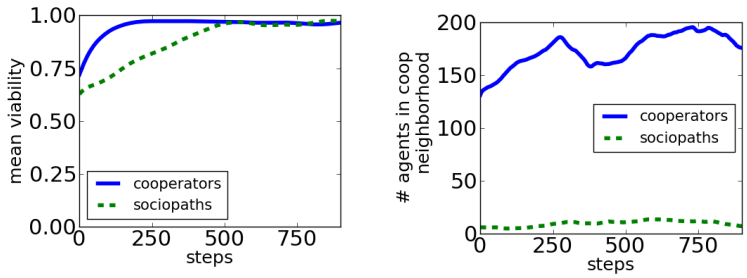

(a) (b)

Figure 7.6: A comparison between the viability dynamics of a society of perceiving agents and the mean neighborhood distribution for the cooperative perceiving agents over 1000 steps. The left panel compares the mean viability for cooperative and sociopathic perceiving agents. The right panel shows the degree of clustering of cooperative agents by showing the number of agents of each agent type in the neighborhood of *cooperative* agents. The conscience for the agent society is initialized according to the simulated Coid distribution, where $\phi = 1.4$.

Discussion

From figure 7.5, we can observe that the majority of cooperators are able to cluster throughout the run. At initialization time ($t = 0$), cooperators move quickly towards each other. After clusters of cooperators have formed (between $t = 0$ and $t = 199$), the spatial configuration is relatively stable. That is, cooperative agents tend to stay near to clusters, while sociopaths show more unstable behavior, moving from cluster to cluster.

We can explain this spatial behavior by negative influence of sociopaths on each other and on the cooperative agents, as specified by the payoff matrix (see table 4.7). Clusters of cooperative are 'invaded' by sociopaths, after which the clusters disperse. This is caused by the fact that the sociopaths influences the received utility of the cooperators within this cluster in a negative way, which causes these cooperators to move to another location (as specified by equation 5.1), driving them apart. After a short period, the dispersed cooperators join another cluster, or form a new one. This spatial cycle of cluster forming continues until the end of the experiment.

From these spatial dynamics, we would expect some oscillations or periodicity. Figure 7.6(b) shows the total number of agents of each type for all cooperators. This figure shows oscillations in the number of cooperators in the neighborhood. These oscillations do not perpetrate to the viability dynamics as shown by figure 7.6(a). As can be seen, after $t = 500$, the mean viability of cooperators converges to the upper viability threshold η, and shows only small oscillations. An explanation for the absence of a clear mean viability cycle is that our measure is distorted due to the fact that both populations optimize their mean viability and remain at this upper threshold. We can conclude that the influence of neighbors does not perpetrate into the mean viability values of cooperators. This is caused by the fact that perceiving cooperators agents are 'resistant' against the influence of the sociopaths.

7.1.3 Affective learner agent

The affective learner model adds the final condition of biological agency (see section 2.2.4) as an extension to the perceiving model. Affective learner agents add a closure of the network of internal variables through the implicit self-reinforcement of viability via skill and arousal choice. In this way, the affective learner agent also adheres to the *individuality* condition (as discussed in section 6.2).

In this experiment, we are interested in the implicit influence of skill via the environment on the mean viability of cooperators. A simplification was made by initializing sociopaths as perceiving agents and cooperators as affective learner agents. Hence, sociopaths seek to optimize their viability, but do not learn from experiences. This means that sociopaths receive payoff exactly according to the payoff matrix (see table 4.7), while cooperators can modify these payoffs through arousal choice. Through this simplification, we can keep the dynamics as transparent as possible, so that we can more effectively study the influence of the quality of expected future affect in cooperators.

We have the following hypotheses for the macro-dynamics of the learning agent model. First, we expect that the mean cooperator viability shows periodic movement. Second, we expect that the mean cooperator skill is reciprocally proportional to the future mean cooperator viability. From this, a histeroidal cycle is expected to arise.

We performed experiments on an agent society consisting of affective learner agents to assess whether the dynamics of the society simulate the periodic movement of agent learning and well-being as specified by a histeroidal cycle (see section 2.1.2). As mentioned above, only cooperators are affective learners. Therefore, for this experiment, we chose to only analyze the dynamics of cooperator agents, and leave out the dynamics perceiving sociopath agents.

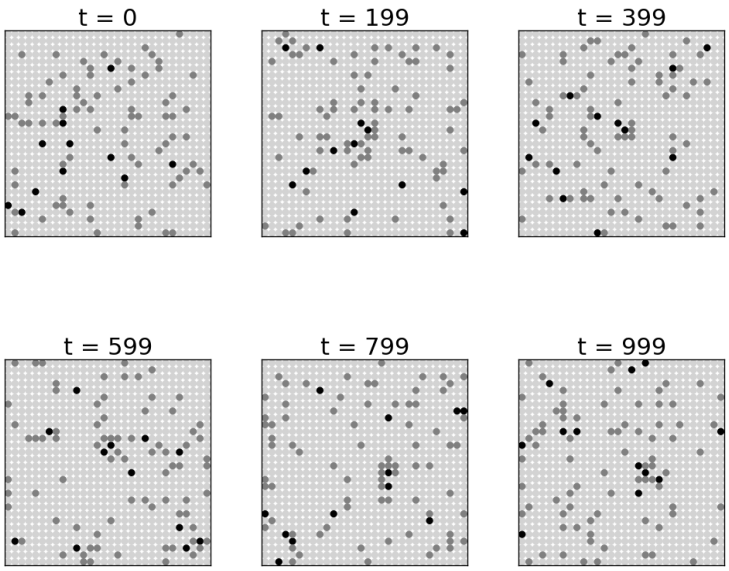

Figure 7.7: Snapshots of the spatial configuration at equally spaced time-steps t for a run of the affective learner agent population over 1000 steps, where the agent society is initialized according to the simulated Coid distribution, where $\phi = 1.4$.

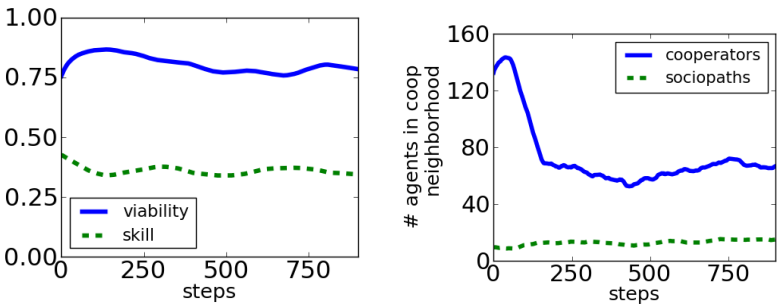

(a) (b)

Figure 7.8: A comparison between the viability dynamics of a society of affective learner agents and the mean neighborhood distribution for the cooperative affective learner agents over 1000 steps. The left panel compares the mean viability for cooperative and sociopathic affective learner agents. The right panel shows the degree of clustering of cooperative agents by showing the number of agents of each agent type in the neighborhood of *cooperative* agents. The conscience for the agent society is initialized according to the simulated Coid distribution, where $\phi = 1.4$.

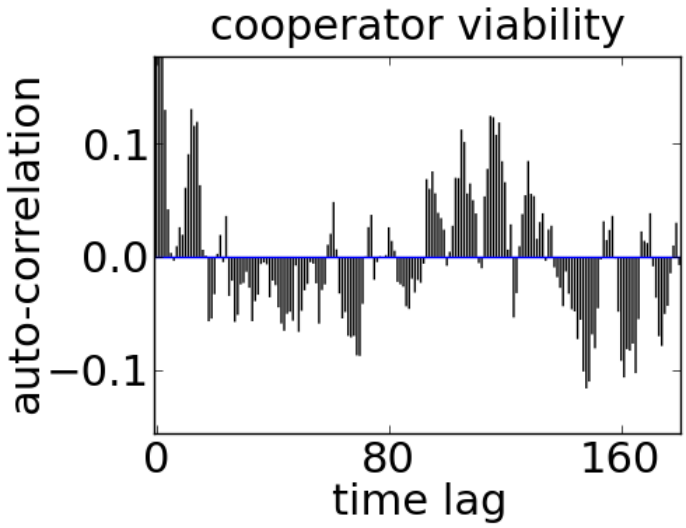

Figure 7.9: A measure of the degree of periodicity present in the mean viability dynamics. The diagram shows the cross-correlation diagram of the mean cooperator viability. The diagram starts at $t = 200$ to ignore the start-up part of the run.

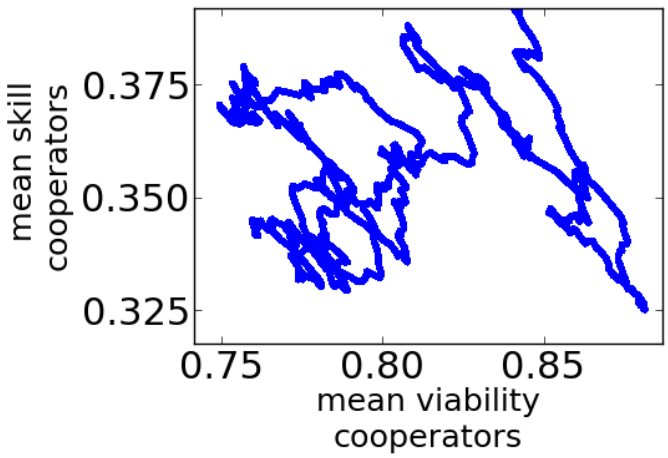

Figure 7.10: A measure of the feedback between mean viability and mean skill over the cooperative affective learner society. The phase diagram shows the relation between the mean cooperator viability at t and mean cooperator viability at $t + 50$. The diagram starts at $t = 200$ to ignore the start-up part of the run.

Discussion: explaining a histeroidal cycle

The spatial dynamics of the cooperative affective learner society (see figure 7.7) are as follows. First, all cooperators gather into clusters, while sociopaths typically gather around these cooperator clusters. After a while, the skill of individual cooperators drops according to the skill decay parameter, since the environment for the cooperators is as expected, an nothing is learned. When this happens, cooperator clusters become unstable and the cooperators disperse (around $t = 200$, see figure 7.8(b)). Cooperators are now learning according to the skill learning parameter, and some clustering among cooperators takes place again. After this, the cycle starts again.

Typically, we see that cooperators cluster around sociopaths. This could be attributed to the fact that sociopaths form a blockade around which cooperators can gather, due to the fact that sociopaths show less noise in their movement: they stick to cooperators if they find them, because they are perceiving agents. These mixed clusters are stable for some time, until too many sociopaths enter the cluster, and cooperators experience a negative viability change, after which they will move elsewhere.

When looking at the individual agent (see figure 7.11(a)), we can observe small cycles of the mean skill and viability values more clearly. The skill drops in times of prosperity (high viability), caused by a decrease of the number of cooperators in the environment (see figure 7.11(b)). Almost immediately, the viability drops, and the skill increases again, after which the viability also increases, provided enough cooperators are present in the environment.

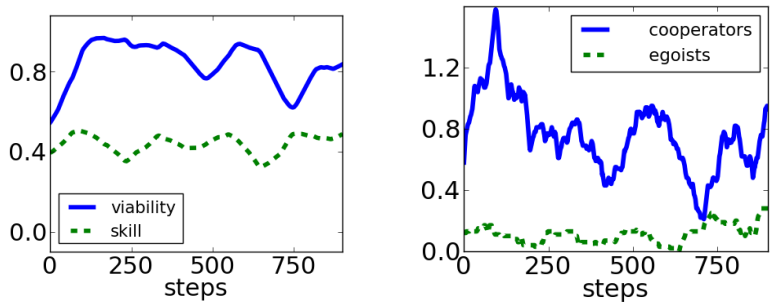

(a) (b)

Figure 7.11: A comparison between the viability dynamics over 1000 steps of a *single* cooperating affective learning agent over in the society displayed in figure 7.7. The left panel shows the viability and skill dynamics of this agent. The right panel shows the degree of clustering of this agent with other cooperative agents by showing the number of agents of each agent type in the neighborhood of this agent.

The learning effect of the skill on the viability is not very apparent for the mean dynamics of the whole population of cooperators, since when skill values increase above the 0.5 threshold, the agent can already optimize its viability. This results only in very small peaks in the viability dynamics. When looking at the macro-dynamics displayed in figure 7.8(a), this learning effect is not seen at

all. This can be explained by assuming that the mean viability cycle, observed in individual agents, is averaged out.

Finally, we observe that the length of the cycles caused by the interplay between skill and viability becomes smaller over time (see figure 7.10). We explain this phenomenon by the convergence of the skill value to the 0.5 threshold for skill learning. At initialization time, the skill value crosses the threshold line less often, hence the cycles are longer. When time continues, the skill value becomes closer and closer to the threshold, which results in a higher number of threshold crossings.

When focusing on the oscillating (histeroidal) nature of the mean viability dynamics of cooperators, figure 7.9 shows oscillations of small positive and negative correlations after $t = 200$. We hypothesize that these oscillations are caused by the learning dynamics of the cooperating agent. By this, we propose that the mean viability dynamics shows some characteristics of a histeroidal cycle.

7.2 Experiment 2: baseline model and extensions with evolution

In this section, we will look at the baseline model and its two extensions in an environment in which agents can reproduce and die in a manner specified by the Demographic Prisoner's Dilemma (DPD) paradigm as described in section 2.2.2. We will compare the macro-dynamics of the extensions to the baseline model *with* evolution, which is analogous to the DPD model. Our model only implements reproduction *without* mutation, as is done in the first two experiments by Epstein [14]. The following rule fully represents the evolution of all agents:

$$\begin{cases} \text{die} & \text{if } v_i < \psi \\ \text{reproduce} & \text{if } v_i > \rho \end{cases} \tag{7.1}$$

Where ψ is the death threshold, ρ is the reproduction threshold. The viability of a newborn agent is initialized according to the viability transfer parameter η (see table 7.1 for the values of the parameters). As can be seen, our model is designed in such a way that it is easily extendable. For instance, we implemented the continuous conscience parameter as the agent 'personality', instead of discrete classes, as in the DPD model. Consequently, the baseline model is analogous, but not equal to the DPD model. Hence we see our baseline model as a re-implementation of the DPD model by Epstein [14]. Table 7.3 lists the agent parameters for this experiment.

symbol	description	value
G	viability change rate	0.008
L	skill learning rate	0.026
C	skill decay rate	0.003

Table 7.3: Agent parameters, equal for all agents. Changes from these values will be mentioned in the text.

We initiate the agent population in a uniform way, as was done in the first experiment in section 7.1.1. We use this uniform initialization throughout this section to test our model with evolution. We chose this distribution so that we can compare the model with the DPD model, which also uses a uniform distribution to initialize the agent conscience. For this experiment we propose that if the full model, the affective learner agent, does replicate the results of some of the extensions of the DPD model, then the affective learner model can be seen as representing the low-level affective motivation of behavior underlying the DPD dynamics.

7.2.1 Baseline agent

This section will test the macro-behavior of the baseline model *with* evolution. As mentioned in section 7.1.1, the baseline model only adheres to one condition of biological agency. In this case, we use the baseline model to both compare it to the its extensions, as well as the DPD model by Epstein [14].

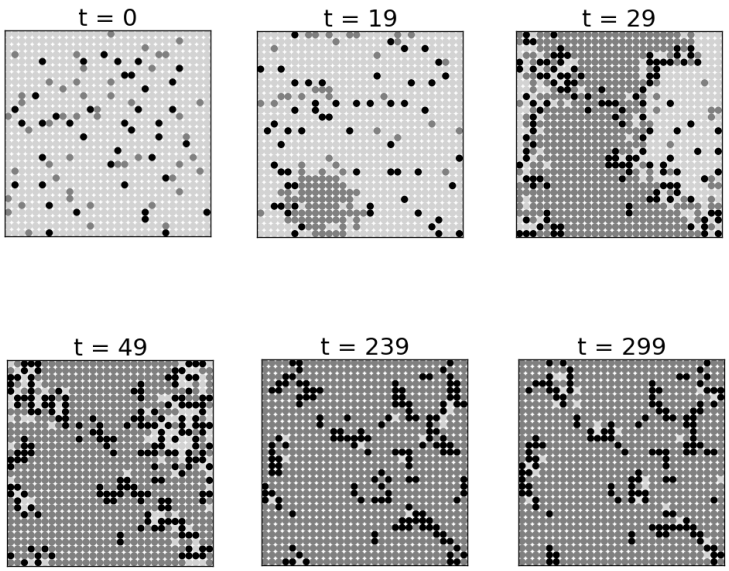

Figure 7.12: Snapshots of the spatial configuration at time-steps t for a run of the baseline agent population over 300 steps with evolution enabled. Here the agent society is initialized according to a uniform distribution.

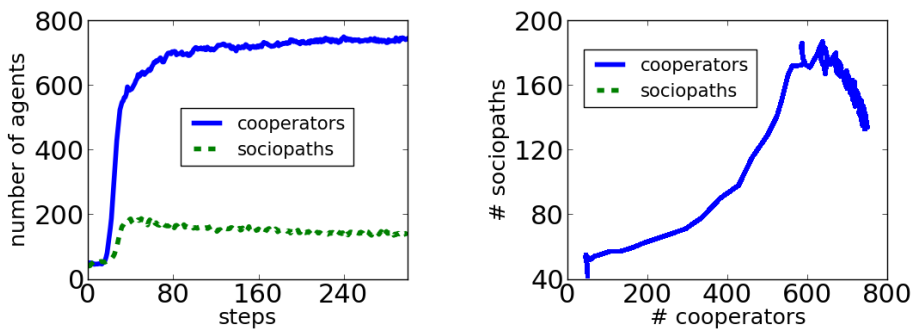

Figure 7.13: Population dynamics for a society of baseline agents *with* evolution over 300 steps. The agents do not have a maximum age. An equilibrium is found around the point (720, 130)

Discussion

We can observe from figures 7.12 and 7.13 (left panel) that the population count for the cooperative agent quickly increases, until the torus is almost filled with agents. This phenomenon can be explained by the fact that cooperators positively influence the received utility of each other. This causes the viability to rise above the reproduction threshold ρ, and a 'clone' of the agent is made into one of the neighboring sites (if the site unoccupied). In this way, neighborhoods of cooperative agents are formed, which is clearly seen in figure 7.12 at $t = 19$

and $t = 29$. Figure 7.13 (right panel) shows that the population dynamics are attracted to a region in the phase space which corresponds to an unstable equilibrium formed around $t = 80$ (as shown by figure 7.12).

The results of this experiment as shown by replicate the results by Epstein in a qualitative way. We see that a dynamic equilibrium arises very quickly. This is due to the formation of cooperative clusters. We can state, following Epstein, that "[...] cooperation can emerge in a population of tag-less agents playing zero-memory fixed strategies of cooperate or defect" ([14], p. 207). Hence, the equilibrium of population dynamics is heavily dependent on the spatial nature of the social space design. Clusters of cooperative agents 'protect' its inner inhabitants from negative affective influences, since sociopaths are not able to travel to the inside of these clusters. The formation of clusters fully determines the population dynamics.

Concluding, we can see our model as a validation of the original DPD model. Furthermore, we can use the baseline model to study the effects of its extensions. The next section will consider the macro-dynamics of the extensions of the model and compare them with the results described in this section.

7.2.2 Perceiving agent

As explained in section 5, the perceiving agent takes deliberate actions through the prediction of future affect. This experiment will test the macro-dynamics of the a society of perceiving agents, when evolution is enabled.

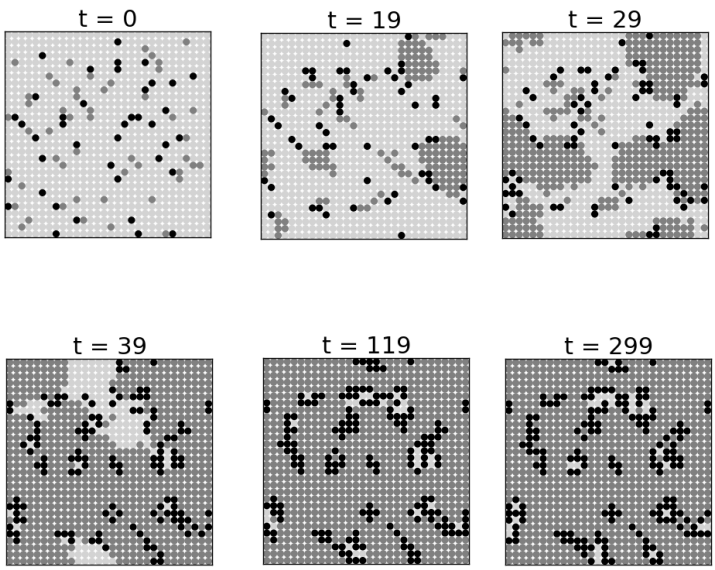

Figure 7.14: Snapshots of the spatial configuration at time-steps t for a run of the perceiving agent population over 300 steps *with* evolution enabled. The agent society is initialized according to a uniform distribution.

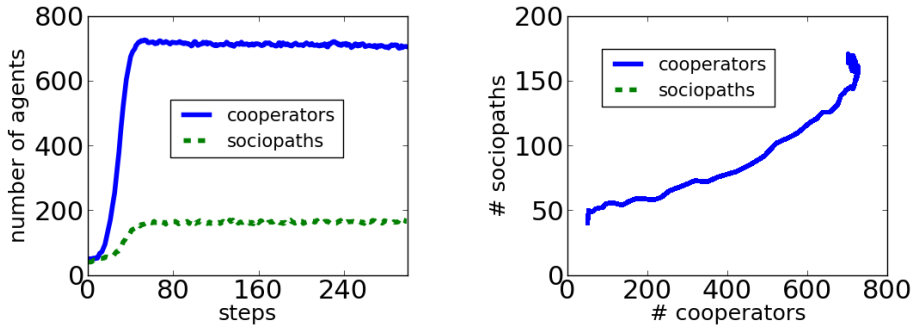

Figure 7.15: The left panel shows the population dynamics for a society of perceiving agents *with* evolution over 300 steps. The right panel shows the phase diagram, representing the ratio between the number of cooperators and sociopaths. An equilibrium is found around the point (700, 150).

Discussion

The population dynamics as depicted by figure 7.14 is qualitatively comparable to the population dynamics of the baseline model and the first two experiments of the DPD model [14]. The cooperators first reproduce quickly. This is illustrated by figure 7.14, where we see the formation of growth of clusters between $t = 0$ and $t = 39$. The cooperative perceiving agents reproduce seem to reproduce more quickly than they do in the baseline model. This can be explained by the inherit drive to clustering of perceiving agents in general, as shown in section 7.1.2. Bear in mind that all agents are perfect in predicting their future affect, hence they always take the optimal action.

Eventually, at around $t = 70$, the population dynamics converge to an equilibrium, as depicted by figure 7.2.2 (left panel). When comparing figures 7.2.2 (right panel) and 7.13 (right panel), we see that the perceiving agent model shows a less noisy equilibrium than the baseline agent model. The baseline agent model shows a small decrease of sociopathic agents, whereas the number of sociopaths does not decrease in the perceiving agent model. Again, this can be explained by the fact that agents are perfect in making decisions bases on the prediction of future affect, hence sociopathic agents are able to survive by always making optimal decisions, by keeping away from neighboring sociopaths, when possible. Concluding, as mentioned, the perceiving agent model is far from being cognitively plausible, since all perceiving agents perceive the environment optimally. In the next section, we will see what the effect of the introduction of learning through arousal choice will have on the macro-behavior of the model.

7.2.3 Affective learner agent

The learning agent adds uncertainty about the future. Agents no longer make optimal decisions based on perfect perception. This experiment will test the macro-dynamics of a society of affective learner agents, when evolution is enabled. Figure 7.2.3 show the results of a run of 300 steps of a society of coop-

erative affective learners and sociopathic perceiving agents, as was explained in section 7.1.3.

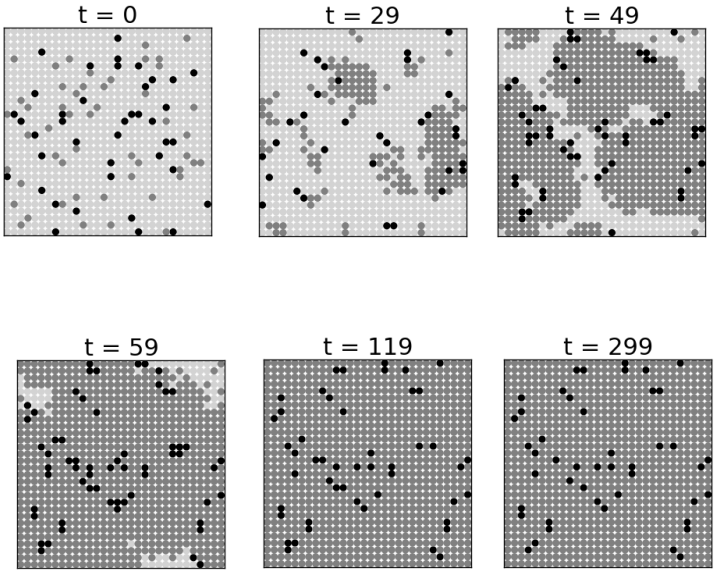

Figure 7.16: Snapshots of the spatial configuration at time-steps t for a run of the affective learner agent population over 300 steps *with* evolution enabled. The agent society is initialized according to a uniform distribution.

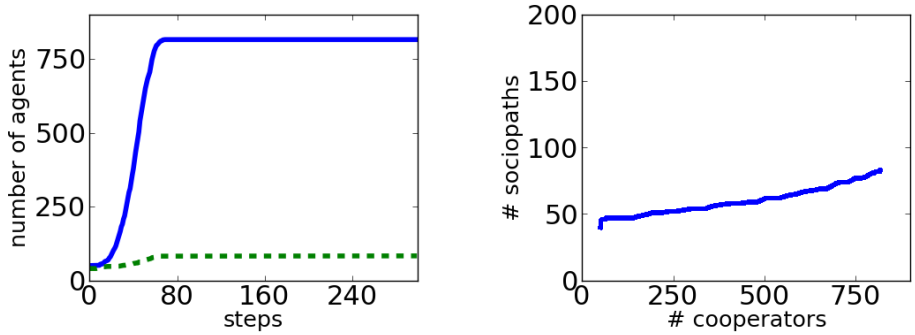

Figure 7.17: The left panel shows the population dynamics for a society of affective learner agents *with* evolution over 300 steps. The right panel shows the phase diagram, representing the ratio between the number of cooperators and sociopaths. An equilibrium is found around the point (780, 70).

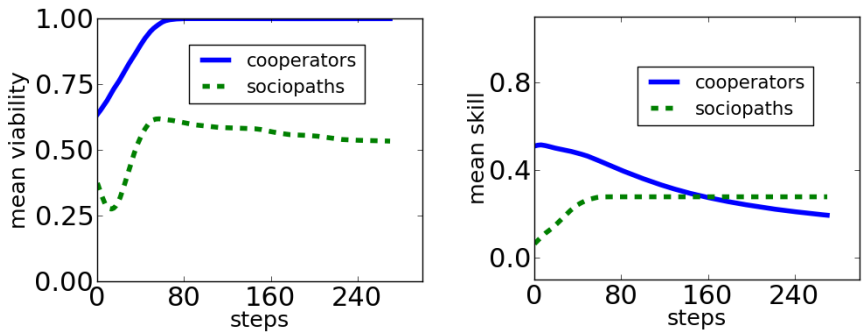

Figure 7.18: The left panel shows the viability dynamics for both cooperators and sociopaths in a society of affective learner agents *with* evolution over 300 steps. The right panel shows the skill dynamics for both cooperative and sociopathic affective learner agents.

Discussion: a histeroidal cycle as basis for population dynamics

In this section, we performed measurements on an affective learner agent society to assess whether the population dynamics of an evolving society can be explained by the periodic movement of agent learning and well-being as specified by a histeroidal cycle (see section 2.1.2). We learned from the experiments in section 7.1 that the skill value has an influence on the viability dynamics of the affective learner agent. From this we would expect that the skill value should have an influence on the population dynamics in an evolving affective learner society, as reproduction and death are intrinsically linked (see equation 7.1.

The results show that affective learner agents and perceiving agents show comparable macro-behavior. The population dynamics, as illustrated by the phase diagram (see figure 7.2.3, right panel) are comparable to that of the perceiving agent (see figure 7.2.2, right panel). First, we see that zones of cooperation form between $t = 29$ and $t = 49$. After this, an equilibrium forms around $t = 80$ (see 7.2.3, left panel), as does happen in the perceiving agent model. The number of sociopathic affective learners is smaller than the the number of sociopathic perceiving agents, at the equilibrium point. We can attribute this to the fact that the mean skill value (as shown by figure 7.2.3, right panel) of sociopaths is low throughout the run. Hence, sociopaths do not always make optimal decisions, and encounter other sociopaths more often than the perceiving sociopaths do encounter each other.

Seemingly, the addition of skill and arousal does not yield oscillating population dynamics, comparable to the viability dynamics of a histeroidal cycle. When we look at the viability dynamics itself, as depicted by figure 7.2.3 (left panel), we are not able to distinguish any oscillating movements. When looking at the mean values, cooperators are able to optimize their viability level, while their skill level drops to 0.2 throughout the run. Sociopaths keep the mean viability level at around 0.5, while the skill level rises to 0.3 throughout the run. For both the cooperative and the sociopathic affective learner agents, we cannot

distinguish oscillations in the mean variable dynamics.

We can attribute the absence of these expected cyclic phenomena in both the internal dynamics and the population dynamics to the quick formation of a large cooperative zone. Within this zone, there is no room to move for the cooperators, hence cooperators are temporarily 'decoupled' from the environment, since they cannot execute their action plans. Only sociopaths can make 'mistakes', caused by a low skill value, by making decisions which have a negative effect on their viability. The absence of a histeroidal cycle can also be attributed to our measurements. It is possible that spatially local periodic cycles are averaged out. All agents only observe a small region around them, hence we would have to introduce a global information variable to see what effects the influence of the sociopaths have on the whole society. Another solution would be to use different local measurements on the data. For instance, we could measure the mean viability for a cluster of agents or a certain local region on the torus.

Chapter 8

Discussion

In this chapter we will proceed to a general discussion. Here, we compare the findings from our experiments with our hypotheses; did the model succeed in answering our research questions? We will also list the shortcomings of our model. After this, we will discuss future applications and extensions. Finally, we end with a general conclusion.

Our model showed the most basic implementation of an affect-based model inspired by a definition of biological agency [11, 2]. We aimed to incorporate the low-level biological idea of an adaptive, self-reproducing system into the domain of social simulation. Our goal was to build a cognitively plausible, affective agent-based model for studying the behavioral foundations underlying the dynamics of co-operation in a society over multiple generations.

We used the Demographic Games approach to build the bridge between the high-level interaction model and the low-level behavior generation model. The affective learner model adds a new dimension to current social simulation research on the evolution of cooperation (Axelrod et al. [1] and Epstein [14]). Through the definition of biological agency, we introduced behavior which is based not in response to the environment, but is based on the influence of the environment upon the system.

We can say the following things about the results of our experiments. First, the DPD social simulation model is not sufficient to be linked to real-world phenomena. We propose that the affective learner model adds cognitive plausibility to the agent-based approach of social simulation, by adding the intrinsic coupling of agent and environment. Second, current PD results can be replicated by the affective learner model: the explanatory power of the model is at least as strong as the DPD model. Third, the generation of the histeroidal cycle shows how the affective learner model provides insight into the relation between the affective state of individuals and the population dynamics of a society. And finally we propose the following:*the affective motivation of behavior is the underlying cause of the dynamics of cooperation.*

We propose that the DPD model is too much a simplification of reality with respect to the assumptions it takes about the origin of behavior, and the spatial nature of a social network. The DPD model does not show why, in the real world, cooperating agents are able to suppress the influence of sociopath agents over

periods in time (through generations). By adding an intrinsic coupling of agent and environment we sought to generate the changing influence of sociopaths on co-operators.

8.1 Shortcomings and future work

The following shortcomings in our model are evident from our experiments. Following the results of our experiments, a number of issues have arisen regarding the current implementation of the affect-based model. We found that the translation from the low-level biological domain to the high-level sociological domain posed a number of problems which could be addressed in future modifications of the model.

First, measures are difficult caused by the dynamic nature of the world. For example, we cannot compare the distance of an agent to a static energy source (see [2]). Furthermore, we did not sweep the individual agent parameter space, hence we could not fully test the robustness of the model. Also, measures over the whole population have the possibility of averaging out local effects.

Second, we implemented the agent action choice as a function of the change of viability. We did this to more easily track the agent movements and decision making. Future work should analyze the macro-effects of taking the actual viability value as criterion for action choices. Third, we did not explore the whole range of spatial and time parameters of the world. It is possible that the model can show stronger effects when it is run for a longer time period. Future work should explore this possibility.

Finally, it is possible that the current dimension of the torus causes some of the effects in our model. For instance, in our experiments, we observed a relatively large influence of blockades of agents on cluster-forming. Future work should also explore larger dimensions.

We have a number of modifications and additions which could serve as the basis for future affect-based models. First, we would like to translate the current discrete implementation to a continuous domain. Second, the skill variable is now implemented in an 'artificial' way; agents exhibiting high skill behave like perceiving agents, while agents exhibiting low skill add the possibility of making detrimental decisions. We would like to replace this system by a real learning mechanism in which agents truly learn through, for example, co-evolution of a viability and skill variables.

Third, two variants of communication between agents can be tested. In the first variant, agents do not know the strategies of other agents; they can only detect the effect of other agents by perceiving changes in their own core affect. In the second variant, agents directly perceive the other agents' strategies, and use a combination of the other agents' strategies and the change of their own core affect to make up a strategy. The agent also does not keep an explicit memory of all encounters. This could be modeled by keeping an explicit history of (all) previous strategy/core affect pairs. An action which is more often associated with a positive core affect has a higher probability of selection.

Fourth, the social network structure could be altered. For instance, non-spatial implementations such as graphs could be used to simulate the social

space of agents. In this way, the (political) structure of social order could also be studied by adding functionality for a hierarchical representation. This way, influence of agents would be tree-like. Agents higher in the hierarchy indirectly reach all agents in the lower layers.

Finally, the agent variables could be made heterogeneous among agents. One may vary the internal parameters among agents, to achieve different personalities of agents. The influence of the interaction range could also be studied. Also, currently, no specific empirical validations of the individual agent model exist. Future work could set up an experiment in which the individual core affect, personality, and social skill of human subjects is measured over time.

8.2 Conclusion

In this research we formalized an agent-based model in which an individual agent adheres to the three conditions of biological agency, while also implementing an interaction model for the application in social simulation. In this way, the so-called *affective learner* agent builds a bridge between the fields of theoretical biology and sociology. We proposed that our novel model can provide new insights into current social simulation research.

Our model is novel in that it utilizes a computational model of emotion as a construct to generate a self-reinforcement mechanism which serves as a simulation of social behavior of organisms in general. Whereas most agent-based models use emotions as a performance element to optimize the behavior of an agent in a particular task, we model affect to gain understanding about the causes of behavioral decisions in the real world.

For our first experiment, we aimed to generate the oscillating dynamics of the mean viability of a society, called a histeroidal cycle. Our experiments showed that a society of agents is able generate the features of a histeroidal cycle. We explained these dynamics as caused by the individual viability dynamics generated by the modulation of the affective impact stemming from interactions between cooperative and sociopathic individuals through internal self-reinforcement. We observed that in our model, the histeroidal cycle only occurs when the agents are equipped with an affective reinforcement mechanism. This mechanism is implemented in the final extension of our individual agent model, the affective learner agent. Furthermore, the histeroidal cycle only occurs in a society in which a majority of cooperating individuals and a minority of sociopathic individuals are present.

For our second experiment, we hypothesized that the histeroidal cycle can be used to gain insights into the low-level behavioral dynamics underpinning population dynamics as modeled by the Demographic Prisoner's Dilemma and other evolutionary game theoretic approaches. Through this experiment, we aimed to generate population dynamics using a society of affective learner agents which have the possibility to reproduce and die. The outcomes of our experiments did not fully confirm our hypotheses. We attribute the absence of the predicted phenomena to the spatial constraints of our world. Clusters of cooperators pose a problem, as they inhibit the self-reinforcing behavior of affective agents. The spatial constraints of the DPD paradigm pose a problem when using more

complex agents. We propose that the conclusions drawn from the generation macro-social dynamics of a society of very simple baseline agents in the DPD paradigm are to be reconsidered, as this paradigm is too much a simplification of reality.

Future work should apply the affective learner model to other agent-based paradigms. The strength of our model lies in its flexibility and inter-disciplinary applicability through its design inspired by insights from three scientific disciplines. First, theoretical biology provides the elementary agent design. Second, insights from emotion research were used to operationalize agent behavior through affective motivation and self-regulation. Finally, game theoretic tools were used to model the interaction design.

Concluding this research, we propose that the affective learner agent offers a novel, cognitively plausible, explanation of the emergence of cooperation among heterogeneous agents in the Demographic Prisoner's Dilemma paradigm. Furthermore, it could provide us with insight into the lower-level foundations of social dynamics. Our approach shows promise in the enlargement of the understanding of macro-social phenomena through the modeling of affective individuals.

Bibliography

[1] R. Axelrod and W.D. Hamilton. The evolution of cooperation. *Science*, 211(4489):1390, 1981.

[2] X.E. Barandiaran, E. Di Paolo, and M. Rohde. Defining agency: individuality, normativity, asymmetry, and spatio-temporality in action. *Adaptive Behavior*, 17(5):367, 2009.

[3] J.H. Barkow, L. Cosmides, and J. Tooby. *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*. Oxford University Press, 1992.

[4] M.E. Bratman. *Intention, Plans and Practical Reason*. CSLI Publications, Stanford, CA, 1999.

[5] R.A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47(1-3):139–159, 1991.

[6] A. Clark. *Being There: Putting Brain, Body, and World Together Again*. The MIT Press, Cambridge, MA, 1998.

[7] J. Coid, M. Yang, S. Ullrich, A. Roberts, and R.D. Hare. Prevalence and correlates of psychopathic traits in the household population of Great Britain. *International Journal of Law and Psychiatry*, 32(2):65–73, 2009.

[8] A.R. Damasio. *Descartes' Error*. Putnam, New York, 1994.

[9] F. de Waal, K. Leimgruber, and A.R. Greenberg. Giving is self-rewarding for monkeys. *Proceedings of the National Academy of Sciences*, 105(36):13685, 2008.

[10] D.C. Dennett. *The Intentional Stance*. The MIT Press, Cambridge, MA, 1989.

[11] E. Di Paolo. Extended life. *Topoi*, 28:9–21, 2009.

[12] E.A. Di Paolo. Autopoiesis, adaptivity, teleology, agency. *Phenomenology and the Cognitive Sciences*, 4(4):429–452, 2005.

[13] M.D. Egbert, X.E. Barandiaran, E.A. Di Paolo, and C.V. Rao. A minimal model of metabolism-based chemotaxis. *PLoS Computational Biology*, 6(12):574–580, 2010.

[14] J.M. Epstein. Zones of cooperation in demographic prisoner's dilemma. *Complexity*, 4(2):36–48, 1998.

[15] J.M. Epstein. *Generative Social Science: Studies in Agent-based Computational Modeling*. Princeton University Press, 2007.

[16] R.H. Frank. *Passions Within Reason: The Strategic Role of the Emotions*. Norton New York, 1988.

[17] N.H. Frijda. *The Emotions*. Cambridge University Press, Cambridge, MA, 1986.

[18] G.N. Gilbert and K.G. Troitzsch. *Simulation for the Social Scientist*. Open University Press, New York, 2005.

[19] R.D. Hare. *Manual for the Revised Psychopathy Checklist*. Multi-Health Systems, Toronto, 1991.

[20] R.D. Hare. *Without Conscience: The Disturbing World of the Psychopaths Among Us*. The Guilford Press, 1999.

[21] J. Hofbauer and K. Sigmund. Evolutionary game dynamics. *Bulletin of the American Mathematical Society*, 40(4):479, 2003.

[22] B.A. Huberman and N.S. Glance. Evolutionary games and computer simulations. *Proceedings of the National Academy of Sciences of the United States of America*, 90(16):7716, 1993.

[23] W. James. What is an emotion? *Mind*, 9(34):188–205, 1884.

[24] C.G. Langton. Studying artificial life with cellular automata. *Physica D: Nonlinear Phenomena*, 22(1-3):120–149, 1986.

[25] J.C. Letelier, J. Soto-Andrade, F. Guíñez Abarzúa, A. Cornish-Bowden, and M. Luz Cárdenas. Organizational invariance and metabolic closure: Analysis in terms of (M, R) systems. *Journal of Theoretical Biology*, 238(4):949–961, 2006.

[26] A.M. Lobaczewski. *Political Ponerology: A Science on the Nature of Evil Adjusted for Political Purposes*. Red Pill Press, Ltd, Grande Prairie, Canada, 2007.

[27] M. Matsumoto and T. Nishimura. Mersenne twister: a 623-dimensionally equidistributed uniform pseudo-random number generator. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 8(1):3–30, 1998.

[28] H.R. Maturana and F.J. Varela. *Autopoiesis and Cognition: The Realization of the Living*. Springer, London, 1980.

[29] John Maynard-Smith. *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, MA, 1982.

[30] J. Nash. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences of the United States of America*, 36:4849, 1950.

[31] M.A. Nowak and R.M. May. Evolutionary games and spatial chaos. *Nature*, 359(6398):826–829, 1992.

[32] M.A. Nowak and R.M. May. The spatial dilemmas of evolution. *International Journal of Bifurcation and Chaos*, 3:35–35, 1993.

[33] A. Ortony, G.C. Clore, and A. Collins. *The Cognitive Structure of Emotions*. Cambridge University Press, New York, 1988.

[34] R.W. Picard. *Affective Computing*. The MIT Press, Cambridge, MA, 2000.

[35] A. Rapoport. Prisoner's dilemma: reflections and recollections. *Simulation & Gaming*, 26(4):489, 1995.

[36] M. Reimer. Psychopathy without (the language of) disorder. *Neuroethics*, 1(3):185–198, 2008.

[37] Don Ross. Game theory. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy (Fall 2010 Edition)*, 2010.

[38] J.A. Russell. Core affect and the psychological construction of emotion. *Psychological Review*, 110(1):145–172, 2003.

[39] K.R. Scherer, E.S. Dan, and A. Flykt. What determines a feeling's position in affective space? A case for appraisal. *Cognition and Emotion*, 20(1):92, 2006.

[40] E. Silverman and J. Bryden. From artificial societies to new social science theory. In M. Capcarrere, A.A. Freitas, P. J. Bentley, C.J. Johnson, and J. Timmis, editors, *8th European conference, ECAL 2005, Canterbury, UK, September 5 - 9*, volume 3630 of *Lecture Notes in Computer Science*, pages 565–574, Heidelberg, 2007. Springer Berlin.

[41] B.R. Steunebrink, M. Dastani, and J.J.C. Meyer. A formal model of emotions: integrating qualitative and quantitative aspects. In *ECAI 2008: Proceedings of the 18th European Conference on Artificial Intelligence*, page 256, Patras, Greece, 2008. IOS Press.

[42] B.M. Stilwell. The consolidation of conscience in adolescence. *Authoritative Communities*, 1:123–135, 2008.

[43] M. Stout. *The Sociopath Next Door: The Ruthless Versus the Rest of Us*. Random House, Inc., 2006.

[44] M. Stout. *The Paranoia Switch*. Farrar Straus & Giroux, New York, NY, 2007.

[45] R.S. Sutton and A.G. Barto. *Reinforcement Learning: An Introduction*. The MIT press, Cambridge, MA, 1998.

[46] F.J. Varela, E. Thompson, and E. Rosch. *The Embodied Mind*. The MIT Press, Cambridge, MA, 1991.

[47] J. von Neumann. *Theory of Self-reproducing Automata*. University of Illinois Press, Champaign, IL, 1966.

[48] J. von Neumann and O. Morgenstern. *The Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1944.

[49] J.W. Weibull. *Evolutionary Game Theory*. The MIT press, Cambridge, MA, 1997.

Appendix A

Graphical user interface

In order to implement the agent-based model, we created a graphical user interface (GUI) to provide ourselves with more insight into the iterative design process. The GUI was implemented in the Python programming language using the wxWidgets library. Figure A.1 shows the main screen for the GUI.

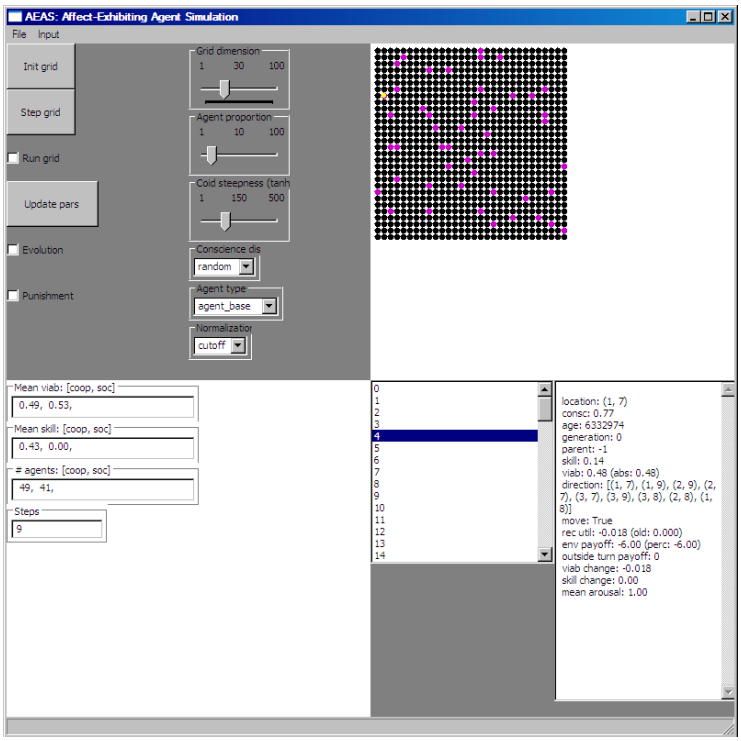

Figure A.1: The main screen of the GUI, which allows the user to track the spatial state of the world, adjust global parameters, and monitor individual agent states.

The upper left panel of the GUI, called the *control panel*, allows the user to start, stop and reset the simulation. Furthermore, it allows for adjusting the size of the field, the number of agents, and the agent extension to use for this simulation. Finally, there are switches to turn on evolution and the experimental punishment feature. The upper right panel, called the *grid panel*, shows a visual representation of the world and the agent types inhabiting the world. The lower left panel, called the *output panel*, shows some of the mean values for some of the internal variables of the agent population, to provide insight into the global internal dynamics of the model. Finally, the lower right panel, called the *agent panel*, shows the internal variables for the selected agent ID. The selected agent is subsequently highlighted by a yellow color in the grid panel.

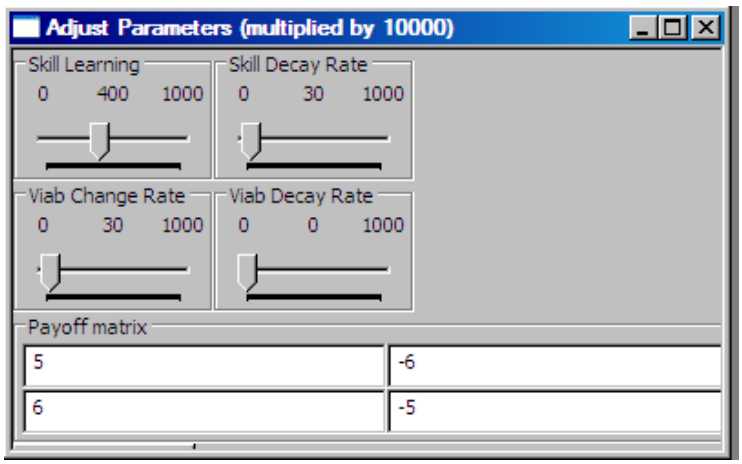

Figure A.2: The parameter adjustment dialog screen of the GUI, which allows the user to adjust the agent parameters on-line.

The GUI aided in the design process in a number of ways. It was primarily designed to speed up, standardize, and simplify the development of the baseline model and its extensions. It provides a quick overview of the model dynamics for the developer, as it shows both individual agent dynamics as well as the global spatial dynamics in which the focussed agent is embedded. The GUI also allows for rapid testing of the behavior of the different agent extensions by adding the possibility of changing parameters on-line (see figure A.2). In addition to the usefulness for the developer, the GUI also makes the dynamics of the model insightful for outsiders. Finally, future developers of the model can more quickly gain an understanding of behavior of the model while developing future extensions.