Universiteit Utrecht

Emergence in Multi-Agent Systems

Technical Artificial Intelligence

MSc Thesis

Author:

Jolanda Duinhoven -
van Steenbergen, BSc

Supervisor:

Dr. Gerard Vreeswijk

Department of Information and Computing Sciences

Faculty of Science

Utrecht University, the Netherlands

June 2012

Abstract

This thesis is meant to clarify the concept of emergence once and for all by creating order from the chaos of emergence definitions. Instead of focussing on the disagreement among scientists on the subject, it will reveal the commonalities between the different ideas on emergence. The resulting unification will create a solid basis for further research on emergence.

Acknowledgements

Contents

Chapter 1

Introduction

The term 'emergence' is believed to be coined by the English philosopher G.H. Lewes in 1875. What started out as a philosophical discussion, eventually spread throughout multiple disciplines, ranging from biology to computer science. The word emergence is derived from the Latin verb *emergere*, meaning 'to come out of the sea' or 'to come to light'. The online Oxford Dictionary defines it as either "the process of becoming visible after being concealed" or "the process of coming into existence or prominence" [21].

The scientific notion of emergence corresponds more to this second interpretation and is related to the well-known principle of the whole being more than just the sum of the parts; a concept already mentioned by Aristotle. Emergence can yield patterns, for instance spatial patterns, or patterns in the form of repeated sequences of behaviour. But emergence can also bring about properties and forms of behaviour that appear to be new.

We encounter emergence in natural systems in the field of biology, chemistry, sociology, and so on. Think for instance of human behaviour, a highly complex phenomenon, ultimately resulting from the 'simple' firing of neurons in our brains. Another example is the distinctive smell of hydrogen sulphide, often associated with the smell of rotten eggs. This is a property that exists at the molecular level, but not at the atomic level; neither hydrogen, nor sulphur (i.e. the atoms that make up hydrogen sulphide) produces this smell when separate. Yet another, and again completely different example, is the emergence of social conventions in human societies, such as shaking hands when meeting someone. For several other examples I refer to ([4], [26]).

Emergence is also observed in artificial systems. In fact, it is said to be

a key ingredient for complex systems, and claimed to lie at the basis of true
artificial intelligence by many. A typical example of emergence in artificial
systems is the 'glider', a moving cyclic pattern occurring in a cellular automata
(CA) called the Game of Life, devised by the British mathematician J.H.
Conway in 1970. It is depicted in figure 1.1. In this CA, every cell in the grid
obeys rules that simply determine the next state of the cell (i.e. dead or alive)
based on the current states of its eight neighbours. A live cell that has less
than two or more than three live neighbours will die from respectively under-
population or overcrowding. A dead cell having exactly three live neighbours
will become alive, as if by reproduction. From these simple rules, complex
patterns (of living cells) evolve that are not explicitly specified by the rules
and cannot be produced by summing over the individual behaviours of cells
when they are isolated from one another. Other examples of emergence in
artificial systems are phenomena like clustering, or convergence to equilibria.
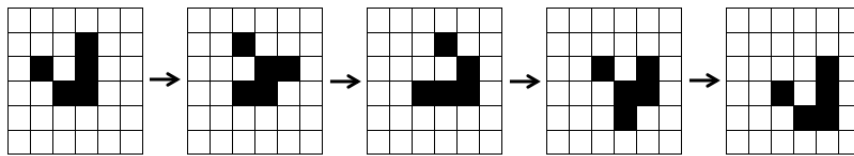
Figure 1.1: Cyclic pattern called 'glider' appearing in the Game of Life.

In conclusion, emergence occurs in many disciplines and does not have a
specific form. This is exactly what makes it so difficult to capture the notion
in a single, explicit and formal definition. It is not surprising then, that it is
common in literature on emergence to construct classifications distinguishing
between different types of emergence instead.

Besides a vast amount of literature on classifications and definitions of
emergence, we can also find numerous works proposing models for the practi-
cal study of emergence in complex systems simulated by multi-agent systems.
See for example ([34], [13], [31]). The main purpose of the study of emergence
in complex systems is to reveal the underlying 'laws of emergence', i.e. the pro-
cesses and properties that constitute emergent phenomena. John H. Holland,
one of the first PhD's in computer science, has paid a significant contribution
to the debate on emergence with his book 'Emergence: from Chaos to Order'
[25] in 1998. One of the things he suggests is that, by studying a wide variety
of emergence examples, we could discover the essential conditions for emer-

gence through elimination of incidental properties. Multi-agent systems are extremely suited for this study since nearly all systems exhibiting emergence can be modelled as multi-agent systems. That is, they are often characterized by multiple autonomous entities interacting with each other.

As soon as we understand what is needed to create emergence, we might be able use this knowledge to design emergence ourselves and use it to our advantage. Today, emergence is already being used as a means of accomplishing desired complex behaviour from simple rules to some extent. But this is still in its infant stage. It is mostly accomplished by means of trial and error.

The knowledge gained from the study of emergence could also be used to find ways to prevent undesirable emergent phenomena (such as deadlocks in manufacturing systems) from happening. Or at least help in finding the cause. A prominent article on this subject is [30], by H. Van Dyke Parunak and R. S. VanderBok.

Furthermore, in simulations of natural processes, gaining insight in emergent phenomena helps us to understand these natural processes. Think for instance of simulations of group behaviour during emergency evacuations, in which we can observe queuing, herding and competitive behaviour. One of several works on agent-based modelling for such phenomena is provided by Wai Kin Victor Chan et al. in [11]. Findings from the analysis of this emergent behaviour might be used by architects to design buildings in which emergency evacuations run more smoothly. The study of emergence in natural processes can even result in finding new solution methods for other systems. For instance, the study of ant foraging behaviour by simulation has brought us algorithms that can find near-optimal solutions to the traveling salesman problem (see [22]).

1.1   Problem statement

After decades of research on emergence, the most pressing question is whether or not the scientific community has established a complete unified theory of emergence. Is emergence fully understood?

However, it turns out that among the many works on emergence, none of them seems to be acknowledged as authoritative or complete by the scientific community. Scientists seem to be nowhere near agreement on even a formal

definition of emergence, let alone on an entire theory of emergence. This division among scientists is also reflected in the works proposing methods for creating, detecting and verifying emergence. It seems very likely that the disagreement concerning a definition of emergence might even be the main reason why a theory of emergence has still not been established.

What I wish to find out within this thesis is what the scientific community *does* agree upon regarding emergence. Exactly how far are we removed from a theory of emergence if we would join several compatible thoughts on emergence?

A full-fledged theory of emergence should at least comprise the following two elements: a formal definition (or classification containing multiple formal definitions) of emergence, and a set of underlying mechanisms of emergence. The first should explain *what* emergence exactly is, and the latter should specify *how* emergence comes about. Within this thesis I will confine my study to these two elements, because they will lie at the basis of a universal theory of emergence.

First, I will set out to examine whether the various definitions and classifications of emergence really are as incompatible as is being implied by the scientists who provide them. Can I extract features that the majority of emergence definitions has in common? And to what extent can I then unify the different emergence classifications?

In the second part of this thesis I wish to establish an inventory of the underlying mechanisms of emergence that have been uncovered by the study of emergence so far. Can I identify the elementary mechanisms of a multi-agent system that are required for emergence?

Currently, there is no known set of minimal requirements that *guarantees* emergence. Even if one would claim to have constructed one, it would be impossible to prove it. Implementations of the same element in different systems are diverse, and it is impossible to examine every possible implementation. For instance, an 'if-then'-rule can take many forms, depending on the domain of the system, and can therefore yield different results. Thus, it is too complicated to construct a complete set of requirements for emergence.

Now what *can* we say about the underlying mechanisms of emergence? Can we at least identify elements of a multi-agent system that increase the chance of emergence occurring in a system, or identify elements that would

preclude the possibility of emergence ocurring in a system? I will attempt to do so within this thesis.

Finally I would like to investigate how these two components (i.e. the classification and the underlying mechanisms) can be used in practice. Can I predict whether emergence will occur based on the mechanisms that can be identified in a system? And can I then also predict which types of emergence might occur?

1.2 Outline

This thesis globally consists of three parts. The first part comprises chapters 2 and 3. In chapter 2, I will treat several classifications of emergence that have received a fair amount of attention from the scientific community. I will compare these different works to one another based on the definitions they provide, revealing quite some similarities. Then, in chapter 3 I will compare these classifications on a more global level, and attempt to construct a joint classification of emergence. The aim of these two chapters is to create order from the apparent chaos of emergence definitions.

The second part of this thesis is formed by chapter 4. In this chapter I will inventory the results of some studies of emergence concerning the underlying mechanisms of emergence. The aim of this chapter is to provide pointers that can be used to identify possible sources of emergence in a multi-agent system. The most important findings of the first and second part are then summarized and joined together in chapter 5.

The third part, chapter 6, comprises a case study of an Artificial Market Model, implemented in NetLogo. I will analyse two observed phenomena that can be considered emergent. In this analysis I will use my findings from the previous chapters. This will reveal to what extent the emergence classification from chapter 3 and the mechanisms addressed in chapter 4 may be suitable for the practical study of emergence in multi-agent systems.

Finally, in chapter 7, I will summarize the most important findings and reflect on how far we are removed from a unified theory of emergence based on the elements I have discussed within this thesis.

Chapter 2

Emergence classifications

In the introduction I have already mentioned that definitions of emergence usually come in the form of a classification of different types of emergence. Unfortunately, it is impossible to study, let alone incorporate, all of the proposed definitions within this thesis. Therefore, I have selected only a few works to discuss here.

Before I proceed to list a selection of such classifications and definitions, I would like to point out that several good overviews of emergence definitions already exist. Some works incorporating such overviews are ([19], [24], [26], [27]). The purpose of this chapter is not to supplement those, but to create a basis for the remainder of this thesis.

Many scientists have come up with their own definitions of emergence, claiming they do not agree with the ones previously defined by others. Also, the fact that there seems to be no unified theory of emergence is often mentioned in literature. In this chapter, I will adopt a different approach. Instead of focussing on the differences, I will point out the commonalities between the various definitions and classifications.

Most classifications of emergence provide rather general definitions. Since I wish to discuss the practical study of emergence in the domain of multi-agent systems later, I will provide translations of the given definitions to this domain if necessary.[1]

The five classifications that will be discussed in the following sections are

---

[1]Although I will try to keep the translations as clean as possible, I am aware of the fact that such translations always bring along the risk of adding or losing meaning due to subjectivity.

not ordered chronologically, but are ordered based on the number of emergence types they distinguish. I will start with the most basic classification of emergence consisting of two types of emergence, and end with the most extensive. I have selected these particular works based on the fact that they have all received a fair amount of attention within the discussion on emergence.

2.1 Chalmers

One of the main ideas in literature is that of a classification of emergent phenomena that distinguishes between 'weak emergence' and 'strong emergence'. This distinction originates from the two main disciplines concerned with the study of emergence. Strong emergence refers to the form of emergence that is most common in philosophical discussions, whereas weak emergence refers to the notion of emergence that is most common in the exact sciences, especially in complex systems theory.

Two works discussing the notions of weak and strong emergence are [2] by Mark A. Bedau, and [10] by David Chalmers, both philosophers. In contrast to Bedau, Chalmers expresses the two notions using the same terminology for both, which makes it more clear. Therefore I have chosen to adopt the definitions of weak and strong emergence provided by Chalmers here.

2.1.1 Strong emergence

Chalmers states that "(...) a high-level phenomenon is strongly emergent with respect to a low-level domain when the high-level phenomenon arises from the low-level domain, but truths concerning that phenomenon are not deducible even in principle from truths in the low-level domain". This means that for strong emergence, new fundamental laws need to be introduced in order to explain the phenomenon.

This type of emergence is often related to the concept of 'strong downward causation'. We speak of downward causation when phenomena at a global, or higher level of organization exert causal influence (e.g. in the form of constraints) on the lower, local level. Now by *strong* downward causation, we refer to the arising of a supervenient downward causal power that is irreducible. It should be note here that, like emergence, downward causation is yet another concept that is hard to define and therefore knows many definitions. Some

discussions on downward causation can be found in ([3], [8], [23]).

Translated to the domain of multi-agent systems, strong emergence refers to the coming into existence of global properties or behaviours through local (inter-) actions at the agent-level, while it is impossible to derive exactly which basic processes have caused it.

2.1.2 Weak emergence

According to Chalmers, a high-level phenomenon "(...) is weakly emergent with respect to a low-level domain when the high-level phenomenon arises from the low-level domain, but truths concerning that phenomenon are unexpected given the principles governing the low-level domain". Regarding the relation between strong and weak emergence, Chalmers claims that cases of strong emergence are likely to be cases of weak emergence too, but cases of weak emergence need not be strongly emergent.

In terms of multi-agent systems, weak emergence refers to the coming into existence of global properties or behaviours that result from the basic rules and interactions put into the system, but are simply unexpected given these rules.

2.1.3 Interpretation

Chalmers explicitly mentions that strong emergence is never deducible to its constituents, but does not address the deducibility of weak emergence. What he does say, is that that in case of weak emergence, the emergent is *unexpected* given the low-level principles. This implies that we cannot predict weak emergence, but it does not necessarily mean that we can not trace back the cause of the emergent after observing it. Intuitively, I would say that weak emergence should be deducible to its constituents. This contrast between weak and strong emergence is also mentioned by Bedau in [2].

Thus, in both weak and strong emergence, higher level complex behaviours or global properties result from low-level processes, but weak emergence can be explained in terms of its constituents, whereas this is not the case for strong emergence.

2.2 Crutchfield

Many definitions of emergence, like the ones given by Chalmers, are rather abstract, and therefore perhaps of less value for the practical study of emergence in multi-agent systems. An alternative classification, that is somewhat more detailed, is one provided by physicist and mathematician James P. Crutchfield in [15]. He explains that "emergence is generally understood to be a process that leads to the appearance of structure not directly described by the defining constraints and instantaneous forces that control a system. Over time 'something new' appears at scales not directly specified by the equations of motion. An emergent feature also cannot be explicitly represented in the initial and boundary conditions. In short, a feature emerges when the underlying system puts some effort into its creation."

Crutchfield describes two types of emergence called 'pattern formation' and 'intrinsic emergence'. His classification is based on the role of the emergent within the system it occurs in.

2.2.1 Pattern formation

Pattern formation results from the interaction of low level simple entities and requires an observer of the system. The resulting patterns often do not have a specific meaning within the system, but they do to the observer. Pattern formation can occur in multi-agent systems in which agents only interact locally (either direct or indirect) and are unaware of what happens at the global level. An example of pattern formation would be a 'glider' in the Game of Life.

2.2.2 Intrinsic emergence

Intrinsic emergence refers to features or patterns that provide additional functionality to the system. Local agent behaviour influences the global state or dynamics *and* vice versa. Agents make individual choices based on the observed global state that emerged from the combined individual choices of all agents in the past. In many models displaying intrinsic emergence, agents do not interact locally, and are thus only connected to each other through the global level. But Crutchfield does not exclude local interactions here.

2.2.3 Interpretation & relation to previous classification

Both pattern formation and intrinsic emergence seem to be subtypes of weak emergence. It has to be noted though, that some systems displaying intrinsic emergence might be explicitly designed to achieve a specific global property through low level behaviour. The design of the global feedback (e.g. the payoff based on the global outcome of the actions of all agents that is used by individual agents to evaluate their strategies) may invoke certain desirable tendencies in behaviour without explicitly controlling the individual behaviour of agents. Think for instance of convergence to global optima that might require sacrifice of individual optima. This would mean that intrinsic emergence is not always unexpected, which does not agree with the definition of weak emergence given by Chalmers. Apart from that, Chalmers does not seem to address designed or intentional emergence at all.

2.3 Boschetti & Gray

Mathematician Fabio Boschetti and modeller Randall Gray [6] base a classification of emergence on the relative causal power of the emergent features that can be generated by the different types of emergence, and add a third type of emergence to the classification of Crutchfield. I have included this work because it seems that the additional type is needed to complete the classification, i.e. it supplements the other two.

2.3.1 Pattern formation & intrinsic emergence

Boschetti and Gray suggest that pattern formation does not imply causal power or control since the patterns can not be manipulated directly, but only by rewriting the original code. In contrast, intrinsic emergence does imply causal control, since theoretically it would be possible to manipulate the low level behaviour by changing the values of the emergent feature. The behaviour of agents in a stock market, for example, could be influenced by changing the actual value of some stock. Boschetti and Gray note that this is undesirable though, since interaction with computation would blur the distinction between algorithm and input data.

2.3.2 Causal emergence

The additional third category, causal emergence, is the arising of structures that can be controlled without manipulating the lower level constituents. Human beings are an example of entities with emergent causal power. For instance, one person can ask another to do something, and then the other will use his own reasoning to consider accepting the invitation. The person requesting something does not have to re-program complicated instructions into the inner workings of the other persons' mind to bring about the desired behaviour. Thus, in this example we see a a macro-level entity (i.e. a human being, built up from micro-level elements such as neurons) communicating with, or manipulating, another macro-level entity.

2.3.3 Interpretation & relation to previous classifications

Causal emergence relates to the notion of 'strong emergence' mentioned before. The only difference that comes to mind, is that it is not explicitly stated that causal emergence necessarily implies that it is not deducible to low-level mechanisms. Of course, in the previous example it is extremely hard to deduce the low-level processes that constitute the behaviour of the person that was asked to do something. But I am not entirely convinced it is impossible. Maybe we will gain enough insight into the workings of our human brains to do so in the (distant) future.

2.4 Cariani

Like Boschetti and Gray, biologist Peter A. Cariani [9] distinguishes between three types of emergence as well. But Cariani's approach is different. He defines a *formally* based type of emergence he calls 'computational emergence', a *physically* based type called 'thermodynamic emergence', and a *functionally* based type called 'emergence-relative-to-a-model'.

2.4.1 Computational emergence

Cariani states that computational emergence "(...) is concerned with what kinds of global behaviors can be built up through the operation of local rules without altering the local rules." He states that this is a bottom-up proces;

there is no feedback from the global to the local level. So, in terms of a multi-agent system, agents (inter-)act only locally, and are unaware of the emergents at the global level resulting from their combined individual actions.

2.4.2 Emergence-relative-to-a-model

By emergence-relative-to-a-model Cariani means the type of emergence "(...) concerned with the formation of global structures which then constrain and alter local interactions". In a multi-agent system, the result of the collective actions of all agents can influence the behaviour of individual agents. For instance, an agent can adopt a different strategy when the current global outcome is not beneficial to him. So, agents influence the global level by their local (inter-)actions, and vice versa, creating a feedback-loop.

2.4.3 Thermodynamic emergence

The definition of the third category of emergence proposed by Cariani, thermodynamic emergence, is not entirely clear to me at this point. One of the ideas Cariani seems to propose, is that each type of emergence results in some form of order, and that the origins of the resulting order are different for each type of emergence. He argues that computational emergence can be associated with order-from-order, emergence-relative-to-a-model can be associated with order-from-chaos, and thermodynamic emergence can be associated with order-from-noise. Regarding thermodynamic emergence, Cariani also mentions the 'emergence of new structures through fluctuations'. Furthermore, 'the origins of life' is listed as an example of thermodynamic emergence.

2.4.4 Interpretation & relation to previous classifications

Cariani's computational emergence can be easily related to pattern formation, whereas emergence-relative-to-a-model clearly relates to intrinsic emergence as defined by Crutchfield. Computational emergence and pattern formation both result from feed-forward relations, while emergence-relative-to-a-model and intrinsic emergence both require feedback from the global level in addition to those feed-forward relations.

Concerning thermodynamic emergence, the fact that 'the origins of life' is listed as an example of this type of emergence seems to imply a relation

between thermodynamic emergence and strong emergence, but it is unclear to me to what extent those two notions are related exactly.

2.5 Fromm

The final work I would like to discuss, is a more detailed classification proposed by physicist and computer scientist Jochen Fromm in [24]. His classification consists of four types of emergence, of which the first three are further subdivided into two types each. Similar to several other classifications, Fromm distinguishes different forms of emergence based on the underlying mechanisms causing it. However, his descriptions are in terms of components of multi-agent systems. Therefore, Fromm's classification of emergence promises to be useful for the practical study of emergence in multi-agent systems.

2.5.1 Type I

Emergence of the first type, called 'simple emergence', can be either intentional or unintentional, and occurs in systems without top-down feedback. Type Ia, simple 'intentional' emergence (also nominal emergence, after [3]) refers to the function of a system as an emergent property of its components, where each component has a fixed role. Typical examples of type Ia emergence occur in clocks, steam engines, or other machines. The behaviour of each part of such a machine is independent of the other parts' states, the global state and the state of the environment, but the function fulfilled by the machine as a whole can not be fulfilled when parts are isolated from one another. So, parts interact but do not influence each others behaviour.

Simple 'unintentional' emergence (type Ib) yields a global property that depends on the relations between the many loosely coupled, disorganized and equal elements of a system. Thermodynamic properties like pressure, volume and temperature are examples of this type of emergence. But also physical properties, such as avalanches and wave-fronts, belong to type Ib emergence. A property is a type Ib emergent when it only applies to the whole and not to the parts.

2.5.2 Type II

Emergence of type II, 'weak emergence', includes top-down feedback from the global (or macroscopic) level to the local (or microscopic) level. It can arise through direct, as well as indirect interactions between agents in a multi-agent system. Fromm continues by distinguishing between two subtypes of type II emergence. He defines type IIa emergence, called 'stable' weak emergence and type IIb emergence, called 'instable' weak emergence.

Stable weak emergence can occur in systems with bottom-up influences and negative top-down feedback from the group or the environment. Negative feedback constrains the actions of agents, and usually contributes to the balance of a system. Examples of stable weak emergence are the foraging behaviour of ant colonies and the emergence of optimal prices of goods in an economy.

Stable weak emergence is unintended but positive. Instable weak emergence on the other hand, "(...) is a form of undesirable, negative emergence through positive feedback, until the bubble explodes or vanishes due to exponential growth." Positive feedback reinforces behaviour of lower level constituents, like the intake of an addictive substance activates the production of substances in the human body that make it ask for more of that addictive substance. The crash of a stock market and the formation of ghettos are other examples of instable weak emergence.

2.5.3 Type III

Emergence of type III, 'multiple emergence', occurs in complex systems with many feedback loops or complex adaptive systems with intelligent agents. Fromm distinguishes between two subtypes of multiple emergence.

Type IIIa occurs in systems in which short-range positive feedback is combined with long-range negative feedback, i.e. so-called activator-inhibitor systems. Examples of emergence in such systems are skin patterns of animals (through activating and inhibiting melanocytes) and stock market rushes.

Type IIIb occurs in adaptive and evolutionary systems and refers to phenomena like evolutionary transitions and sudden scientific or mental revolutions.

2.5.4 Type IV

Finally, type IV emergence, or 'strong emergence', can be defined as "(...) the appearance of emergent structures on higher levels of organization or complexity which possess truly new properties that cannot be reduced, even in principle, to the cumulative effect of the properties and laws of the basic parts and elementary components." In this type of emergence, the macroscopic level and microscopic level are separated by a mesoscopic or intermediate level, ensuring independence between the two.

2.5.5 Alternative classification criteria

Fromm argues that his classification remains the same when based on several other criteria. The tables below (copied from [24]) show the specification of the four emergence types in terms of roles of the components in the system, levels of prediction, boundaries, feedbacks and jumps or leaps.

Type	Name	Roles	Frequency	Predictability	System
I	Nominal or Intentional	fixed	abundant	predictable	closed, with passive entities
II	Weak	flexible	frequent	predictable in principle	open, with active entities
III	Multiple	fluctuating	common - unusual	not predictable (or chaotic)	open, with multiple levels
IV	Strong	new world of roles	rare	not predictable in principle	new or many systems

Table 2.1: Table copied from [24]; classification of emergence based on roles, frequency, predictability and system specifications.

Fromm suggests that emergence of something is possible at a clear boundary of a system, and that it is usually associated with a jump or leap to a new level. What I understand is that 'boundary' refers to the border between different levels of a system, or different phases (in case of type IIIb). Fromm has added the columns 'frequency' and 'system' to table 2.1. The first indicates the frequency of occurrence of each type of emergence in general, i.e. how likely it is to encounter this type of emergence. The latter specifies some other basic characteristics of the system each type of emergence may occur in.

Type		Boundaries	Feedbacks	Jump/Leap
I	Ia	agent-system boundary (only in one direction)	no feedback but absolute commands or constraints	intended or static jump to higher level of organization
	Ib	agent-agent boundary	scale-preserving (peer-to-peer) feedback	fluctuations, no jump to significant higher level of organization
II	II	agent-group or agent system boundary (in both directions)	scale-crossing (top-down) feedback, positive *or* negative	dynamic jump to higher level of organization
III	IIIa	agent-group or agent system boundary (in both directions)	scale-crossing (top-down) feedback, positive *and* negative	dynamic jump to higher level of organization
	IIIb	large fitness barriers in complex evolutionary systems	multiple feedbacks in a system	quantum leap in complex adaptive system
IV	IV	boundary between different evolutionary systems, *barrier of relevance*	all of the above, incl. feedback between different systems	gateway or quantum leap in evolution to new (evolutionary) system

Table 2.2: Table copied from [24]; classification of emergence based on boundaries, feedbacks and jumps/leaps.

2.5.6 Interpretation & relation to previous classifications

Obviously, type I emergence is always deducible, so we can rule out a relation between type I emergence and strong emergence as defined by Chalmers. But since type I emergence is not unexpected by definition, we can not completely match type I emergence with Chalmers' weak emergence either.

Because of the lack of top-down feedback, type I emergence seems to correspond with Crutchfield's pattern formation rather than intrinsic emergence. Entities interact locally and are unaware of the global level. The only difference seems to be that an emergent of type I often has a specific meaning within the system, whereas Crutchfield specifies that pattern formation often does not. But since Crutchfield does not exclude this possibility, I believe it is justified to relate type I emergence to pattern formation here.

Finally, type I emergence can easily be related to Cariani's computational emergence since both concern global behaviours resulting from local rules without altering these local rules.

Type II emergence is called weak emergence by Fromm, and relates to the identically named notion provided by Chalmers. However, Fromm's weak

emergence seems to be a subset of Chalmers' weak emergence, rather than a complete match. Unlike Fromm, Chalmers does not specify that weak emergence requires a specific form of feedback. This implies that weak emergence as defined by Chalmers can also occur in systems with feed-forward relations only, whereas weak emergence as defined by Fromm cannot occur in such systems.

Fromm's type II emergence can be related to Cariani's emergence-relative-to-a-model in a straightforward manner since both require top-down feedback and interactions between agents. The same goes for type II emergence compared to Crutchfield's intrinsic emergence. The only difference is that Fromm specifies further subtypes of this type of emergence by distinguishing between different sorts of top-down feedback (i.e. negative and positive feedback). Fromm claims that these different forms of feedback yield different global behaviour.

Type IIIa emergence seems to relate to emergence-relative-to-a-model and intrinsic emergence in the same way as type II emergence does. Again, the main idea is that it requires top-down feedback.

According to Fromm's description, emergence types IIIb and IV seem to be closely related. Both types occur in evolutionary systems and are characterized by a so-called quantum leap (see last column of table 2.2). Moreover, both type IIIb and type IV seem to relate to the definition of strong emergence provided by Chalmers. Although Fromm did not specify that type IIIb emergence is non-deducible to its constituents, the term 'quantum leap' (see last column of table 2.2) seems to imply that this is the case.

Chapter 3

Comparison of classifications

In the previous chapter, I have already pointed out some commonalities between the definitions proposed by various scientists. In this chapter, I would like to discuss a few points concerning the more structural level of the classifications listed before. Section 3.1 addresses the cardinality of emergence classifications, i.e. the number of different categories of emergence that can be distinguished. This will be followed by a discussion on how the different types of emergence relate to one another within a classification.

The aim of this chapter is to show that the different classifications of emergence turn out to be not that incompatible after all. Therefore, the section concluding this chapter, section 3.3, will comprise an attempt to join the compatible thoughts on emergence in one classification. The resulting unified classification of emergence will then be used for the remainder of this thesis.

3.1   Global structure

While writing the previous chapter, I noticed that it seems to be rather common in literature to divide emergence into three main categories. This is the case in the classification proposed by Boschetti and Gray, as well as in the classification provided by Cariani. Furthermore, the two notions of weak and strong emergence are supplemented with a third category, called 'nominal' emergence, by Bedau (see [3]).

When we take another look at Fromm's classification of emergence, we can discover three main categories as well. In section 2.6 I already mentioned

that Fromm's emergence types IIIb and IV seem to be closely related. Given Fromm's description and table 2.2, it furthermore seems that emergence type IIIa relates more to type II than it does to type IIIb. Both emergence types II and IIIa occur at the boundary between the agent-level and the global-level (i.e. the group or system), relations between the two levels work in both directions and it involves some form of top-down feedback. Hence, it seems that type IIIa emergence can be merged with type II and type IIIb emergence can be merged with type IV. This reveals a global structure consisting of three main categories, which can be mapped to the classifications listed previously. I will elaborate on this in the following sections.

Apart from the ones discussed in this chapter, I have encountered several alternative classifications in literature that divide emergence into three categories as well. See for instance [18] by Terrence W. Deacon, [33] by Stephen Jones, and [17] by Hans de Haan. Thus, it seems that many people agree that three main categories are sufficient to cover the entire range of emergent phenomena.

3.2 Hierarchy & complexity

The next point I would like to address, is that the classifications of emergence treated in this chapter all seem to be hierarchical in some form. The choice of the names for the notions of 'strong' emergence and 'weak' emergence is a rather obvious indicator. The table below (table 3.1) was constructed to illustrate an ordering of emergence types by increasing complexity.

high	strong emergence		causal emergence	thermodynamic emergence	type IV	type IV
					type IIIb	type IIIb
	weak emergence	intrinsic emergence	intrinsic emergence	emergence-relative-to-a-model	type IIIa	type IIIa
					type II	type II
low		pattern formation	pattern formation	computational emergence	type I	type I
complexity	Chalmers	Crutchfield	Boschetti & Gray	Cariani	Fromm	Fromm*

Table 3.1: Classifications & complexity

By the complexity of an emergence type I mean the complexity of the underlying mechanism producing this type. This is related to the complexity of the system in which a certain type of emergence may occur, e.g. the amount of different forms of relations between elements and levels within the system. For instance, computational emergence is less complex than emergence-relative-to-a-model since the first only requires local interactions whereas the latter also requires feedback from the global level.

In the column representing Fromm's original classification, I indicated that the two subtypes of type III emergence are of a different complexity. This distinction is based on the fact that type IIIa involves two types of top-down feedback, while type IIIb involves multiple feedbacks. Furthermore, type IIIb occurs in adaptive and evolutionary systems, while type IIIa can also occur in less complex systems.

The column in table 3.1 labelled Fromm* reflects the reorganization of Fromm's classification as proposed in section 3.1. Although I argued that type IIIa belongs to type II emergence, the difference in complexity between type IIIa and types IIa and IIb remains. Subtypes IIa and IIb are of the same complexity, since they both involve only one type of feedback, but type IIc can be considered to be more complex since it involves two types of feedback. The same idea holds for type IIIb emergence, which was merged with the original type IV emergence. Since the first occurs within a single system and the latter occurs between multiple systems, the latter can be considered to be more complex.

The complexity-based hierarchy laid out here poses certain implications for the relations between different types of emergence within the same classification. It makes sense that a system producing highly complex behaviour is also capable of producing less complex behaviour, but that the reverse does not necessarily hold. In section 2.1.2 I have already mentioned that, according to Chalmers, cases of strong emergence are likely to be cases of weak emergence too, but cases of weak emergence need not be strongly emergent. This type of connection also holds for all other classifications. For instance, when we encounter intrinsic emergence in some complex system, then we could also encounter pattern formation in that same system. The elements enabling pattern formation (i.e. local interactions) are a subset of the elements enabling intrinsic emergence (i.e. local interactions & feedback from the global level).

Thus, presence of intrinsic emergence guarantees the presence of the elements that can yield pattern formation.

3.3   Unified classification

In this section I will construct a classification of emergence by joining the classifications that have been shown to be compatible in the previous sections. In section 3.1 I have argued that many classifications consist of three categories of emergence. The main ideas of these three categories of emergence can be summarized as follows:

  I deducible emergence resulting from local feedback only

 II deducible emergence resulting from local feedback
    in combination with feedback from the global level

III 'non-deducible' emergence resulting from local feedback
    and multiple feedback-loops (that can not be pinpointed)

The first category captures Crutchfield's pattern formation, Cariani's computational emergence, and Fromm's type I emergence. It occurs in systems in which agents have only local knowledge, meaning their actions are driven by observations of changes in their local environment (i.e. local feedback) only.

This type of emergence can manifest itself in the form of spatial patterns, possibly with a temporal aspect (e.g. cyclic patterns). But it can for instance also be in the form of solutions achieved through cooperation between agents on a local level, that can not be achieved (or require more time) without cooperation. These emergents do not change agent behaviour and are observed from outside the system. Thus, it depends on the external observer which macro-level phenomena are noticed and considered to be emergent.

In general, the phenomena that are labeled emergent by an observer are phenomena that exhibit some kind of regularity or tendency, are recurrent (i.e. they occur more than once), and can not be predicted based on simple inspection of the elementary rules of the system. The latter corresponds to Chalmers' claim that weak emergents are 'unexpected' given the principles governing the low-level domain. It furthermore relates to the idea presented by Crutchfield that an emergent structure should not be "(...) directly described by the defining constraints and instantaneous forces that control a system."

The second category captures Crutchfield's intrinsic emergence, Cariani's emergence-relative-to-a-model, and Fromm's type II and type IIIa emergence. Local (individual) actions indirectly influence the global level, and the global level influences the actions on the local level. The resulting macro-level emergents play a role in the system in that they influence the micro-level. This can yield coordinated behaviour, for instance in the form of systems converging to equilibria.

In systems displaying this type of emergence, agents have partial knowledge about the macro-level, acquired through feedback from (and about) the macro-level. This feedback often has the form of (potentially negative) payoffs, which reflect, but do not explicitly reveal, the macro-level phenomena resulting from local (inter-) action. The behavioural rules of agents enable the agents to react to the feedback they receive, as well as to possible individual observations. An agent could for instance adapt his behaviour by switching between several strategies that are available to him.

Finally, the third category represents emergence defined by Chalmers and Fromm as strong emergence, which corresponds to Boschetti and Gray's causal emergence, and seems to match Cariani's thermodynamic emergence. For this type of emergence, we need agents that are macro-level entities (i.e. entities composed of elementary entities) that are adaptive and aware of the emergence they cause by their interactions. Such agents should be able to change their behaviour, possibly changing their entire set of behavioural rules, in reaction to the feedback from the global level (and potentially their individual observations). An agent should thus be able to control the lower-level constituents it consists of. This results in circular causality. Furthermore, the example of communication between macro-level entities given in the description of causal emergence in 2.3.2, seems to coincide with the other definitions adressed here as well. Thus, a macro-level entity can influence the behaviour of another macro-level entity without directly manipulating its lower level constituents.

One thing that has to be noted here, is that scientists do not agree on whether or not emergence of this third category is strictly non-deducible. Some claim that everything is eventually deducible, provided we have sufficient knowledge of the underlying mechanisms. However, the deducibility-discussion goes beyond the scope of this section.

For the remainder of this thesis, I will adopt the terms 'category I emergence', 'category II emergence' and 'category III emergence' to refer to the classification described in this section, that resulted from joining the works discussed in chapter 2.

Chapter 4

Mechanisms of emergence

The aim of this chapter is to indentify the underlying mechanisms of emergence that have been revealed so far. As I have mentioned in the introduction, emergence is mostly studied in multi-agent systems, since nearly all systems exhibiting emergence can be modelled as multi-agent systems. A multi-agent system is an organization defined by an environment in which multiple agents (inter-) act by applying rules. For each of these components of multi-agent systems, I will list several elements that are required for, or potentially enable emergence. These elements are derived mainly from the findings of John H. Holland in [25], and my personal observations of examples of sytems exhibiting emergence.

Another source I will use in this chapter is [26], written by Aleš Kubík. This work has recieved a fair amount of attention by the scientific community studying emergence and provides some additional insights to the work by Holland.

The Game of Life is a typical example of a system displaying emergence. Scientists generally agree that the Game of Life produces certain patterns that are considered emergent. It is probably the most studied example, and will therefore also appear frequently throughout this chapter.

Some of the elements that will be treated may seem rather trivial, but they cannot be left out because I aim to provide an inventory that is as complete as possible.

4.1 Environment

The environment of a multi-agent system is a structured entity shared by all agents. It defines or visualizes the (potential) connections between agents.

The structure of an environment can be either spatial, organizational, or both. In chapter 6 we will see an example of traders in an artificial financial market, where each trader is visualized as a cell in a two-dimensional grid. The purpose of this spatial representation is to visualize the close colleagues (i.e. the eight surrounding cells) potentially influencing an agents' decisions. Of course, the environment can also simply be the actual space in which agents move around.

Organizational structures are present in systems implementing agents with different roles, defining relations between those roles (e.g. hierarchy). For the remainder of this chapter I will mainly consider spatial structures, because they provide a visual representation that is less abstract.

4.1.1 Indirect interaction

Agents can also interact with each other through the environment, by manipulating (objects in) the environment. Agents can for instance pick-up, move and place objects in the environment.

An example of a system incorporating this form of indirect interaction between agents is an ant colony, in which ants leave pheromone trails for other ants to follow. In this particular system, trails indicating shortest paths between the nest and sources of food emerge. Although these emergent structures are beneficial to them, the ants themselves are not aware that the paths they form and follow are shortest paths. Thus, this example shows that direct interactions are not necessary for producing category I emergence. Systems incorporating indirect interactions *can* produce category I emergence too.

Kubík [26] claims that interactions through the environment can be a source of emergence, but that it is not sufficient for emergence without some additional requirements. According to Kubík, an example of such an additional requirement is that parallel actions of agents should not be replaceable by sequential actions without altering the resulting patterns. I will elaborate on this in section 4.3.4.

4.1.2 Environment role

The environment can also play an active role in a multi-agent system. It can have its own properties and rules and it can provide feedback to agents. An example of an environment property is the local temperature. We could for instance have a system in which each agent prefers to live in an area with a certain temperature, and that an agent is capable of creating fire. When an agent lights a fire, the environment will provide feedback in the form of an increase in temperature.

An environment playing an active role in a system could be regarded as an agent, i.e. a certain type of autonomous entity that knows just one instantiation. This entity is connected to each agent, thus contributing to the level of connectivity within the system. I will elaborate on the level of connectivity in section 4.2.1.

4.2 Agents

Agents in a multi-agent system are autonomous entities, i.e. they sense and act according to their own rule set. Besides a rule set, agents can possess certain properties defining their state or abilities, e.g. the colour of an agent, the amount of money an agent carries, or the distance an agent can travel in one time step. In case of a spatial environment, agents are usually also specified by their current location.

4.2.1 Connectivity & dependence

The level of connectivity between agents within a system seems to be an important factor for emergence. Cellular automata such as the Game of Life, have an obvious network-like structure. Agents are connected to neighbours in adjacent cells and are *dependent* on them, i.e. the states of an agents' neighbours determine his future state.

In the Game of Life, each agent has eight neighbours, i.e. the agents inhabiting the cells that are horizontally, vertically and diagonally adjacent to his. But we can also define smaller or larger sets of neighbours. Emergence results from interactions, and the higher the level of connectivity, the more interactions take place. Thus, the amount of others an agent is dependent on,

might influence the scale of the emergent phenomena that occur in a system.

In chapters 2 and 3 I have already mentioned that behaviour and structures are generally recognized as emergent when they exhibit some kind of regularity or order. When agents are highly connected (i.e. connected to many others), this could lead to more unity among agents than when connectivity is low. And intuitively, more unity implies more regularity, hence an increased chance of emergence.

In other systems we can detect network-like structures as well. Take for instance crowd simulation. At time t, an agent is connected to the other agents that he perceives at time t for the movement of those other agents determines the movement of the agent. In systems like these, in which agents are mobile, connections between agents are dynamic.

4.2.2 Quantity

The population size of agents in a multi-agent system can be a requirement for emergence. Obviously, all emergent phenomena require at least two agents, but this is already guaranteed when we look at *multi*-agent systems. What I actually mean is that there are emergent phenomena that require a specific minimal amount of agents. For example, consider an emergent pattern consisting of n agents. If there are less than n agents present, this pattern cannot be formed.

The population size of agents in a multi-agent system can be variable. In some systems, agents can die, replicate, reproduce, etcetera, but there always is at least one point in time where the population size is at least two. Systems in which the population size fluctuates can entail specific emergent phenomena, such as patterns in fluctuations of population sizes, or tendencies in sizes of populations of certain types of agents. For example, in a predator-prey model with two predator species, one predator species could suddenly outnumber the other species.

It seems rather trivial that the larger the global population, the more interactions take place, thus the more chance that emergence will occur in the system.

4.2.3 Heterogeneity

Variety in rules, properties and roles of agents can increase the chance of emergence, since it would increase the complexity of a system. It seems rather intuitive that the larger the configuration-space (i.e. the number of possible configurations of the system) and the larger the number of different interactions, the more possibilities for emergent macro-level phenomena to arise.

However, heterogeneity in behavioural rules is not *necessary* for emergence. At least not for category I emergence. This is illustrated by the fact that emergent patterns also arise in the Game of Life, in which each agent acts according to the same rules. There is only one heterogenous element required in this system; in the initial situation, there should be heterogeneity among agent states. In the Game of Life, if all agents have the same state in the initial situation, nothing will change.

We could possibly circumvent this by adding some extra rules to the rule set of each agent, specifying that a live agent with eight live neighbours will die and that a dead agent with eight dead neighbours will become alive. This will drastically influence the macro-level behaviour, which could mean that the system no longer produces emergent patterns. I have not tested this, so it is possible that the resulting, completely homogeneous system, still displays category I emergence.

4.2.4 Decentralization

We speak of decentralized (or distributed) knowledge when agents do not have complete knowledge of the system and events ocurring in the system. Recall that in 3.3 I mentioned that category I and II emergence occurs in systems in which agents only act on local observations and the feedback they receive, and have incomplete knowledge of the global level. When each agent would have complete knowledge, it would be easier to produce and explain coordinated behaviour, which would therefore possibly no longer be considered emergent.

By decentralized control we mean that there is no cental leader, coordinating and completely controlling the actions of all (other) agents in order to fulfill a central goal. Centralized control would exclude the possibility of emergence ocurring in a system. If every action of an agent is dictated by a central organization, we will be able to derive all phenomena directly from the

rules of the system. And in 3.3 I claimed that phenomena that are a direct result of executing elementary rules of a system are considered non-emergent.

However, decentralized knowledge and control are inherent to multi-agent systems, so these elements will not exclude the possibility of emergence.

4.2.5 Adaptivity

In evolutionary multi-agent systems and systems incorporating agents capable of learning, the behavioural rules, as well as agent properties (e.g. an agents' abilities), may change over time. Agents adapt their behaviour to the feedback they receive or to their own observations, creating complex feedback-loops. In section 3.3 I have determined that this is an ingredient of category II emergence.

The fact that the behavioural rules of an agent may change, makes it more difficult to deduce observed macro-level phenomena to their micro-level constituents. Tracing back the cause of a macro-level phenomenon requires information about the set of behavioural rules of each agent (involved in the coming about of the phenomenon) at each moment in time. Since keeping this information requires a lot of memory, at some point we will not be able to deduce macro-level phenomena due to limited computational power.

Again, I seem to have stumbled upon the deducibility-discussion. I will not go into the details of this discussion here, but it does point us towards category III emergence.

4.3 Behaviour

Agents can interact with other agents and / or the environment through execution of behavioural rules. A behavioural rule usually has the following form:

$$condition(t) := \sigma(t) \longrightarrow \sigma'(t+1)$$

If the condition is fulfilled at time t, the rule will be applied to the current situation $\sigma(t)$, resulting in a changed situation in the successive time step (indicated by $\sigma'(t+1)$). The condition of a rule specifies things such as states and locations of nearby agents and objects, or internal states of the agent. The current situation, $\sigma(t)$, describes the states of the elements that will be

changed after executing the rule. For example, in the Game of Life, a live agent will die when it has less than two live neighbours. This can be formulated as a rule in the form defined above. The condition states that for this rule to be applied, an agent should be alive and have less than two live neighbours. The current situation contains the fact that the current state of the agent is 'alive'. And the resulting situation contains the fact that at time $t + 1$ the state of the agent will be 'dead'.

The current situation is often part of the condition, but this does not necessarily need to be the case. For instance, we could add a rule to the Game of Life saying: if the neighbour to your left is alive, change your state to 'happy', no matter what your current state is.

For the remainder of this section, I assume that every rule specifies at least one change. I ignore the frame problem by assuming that every change is explicitly specified by the rules and that there is no need for specifying things that remain the same. For example, we can leave out the rule in the Game of Life saying that a live agent will stay alive when he has two or three live neighbours. This is already guaranteed by the fact that in such a situation, the agent cannot apply the rule saying he will die, which requires less than two, or more than three live neighbours.

4.3.1 Chain reactions

Holland has developed a model for formally describing multi-agent systems displaying emergence in [25]. In this model he calls *constrained generating procedures* (*cgp's*), Holland defines a mechanism to be something which has an internal state and determines its next state through a function of the current state and a set of inputs. Each agent in a multi-agent system, defined by a state and a set of transition functions, is an elementary mechanism of the system. These elementary mechanisms are micro-level ingredients of a system and can be linked when the output of one mechanism can be used as input for another mechanism.

In systems displaying emergence, an emergent feature is the product of a series of linked elementary mechanisms. This series of linked mechanisms is simply considered to be a mechanism on a different level, i.e. a macro-level mechanism. Like an elementary mechanism, a macro-level mechanism takes an input, has an internal state (i.e. the combined states of its component

mechanisms), and determines its next state based on the input and its current state.

Recall the 'glider' occurring in the Game of Life (see figure 1.1 in the introduction). Each cell is an elementary mechanism that is linked to its eight neighbours. A cell uses the output of its neighbours (i.e. their states) as input for its 'decision' whether or not to change its state. The resulting state is its output that will be used by its neighbours as partial input for *their* decisions. The glider can be considered a composite mechanism, consisting of nine cells, that takes input from the sixteen surrounding cells which, together with its internal state (i.e. its current configuration), determines its next state. If its sixteen surrounding cells are dead, it will proceed to the next state in his cycle.

An important aspect of the chain reactions depicted by these mechanisms is that they comprise some form of reciprocity or cyclicity. This corresponds to the idea of feedback, which was incorporated in the descriptions of the different types of emergence provided in section 3.3. In the Game of Life we have that the state of a cell partially determines the state of its neighbouring cells *and* vice versa, i.e. the state of a neighbouring cell partially determines *its* state. If we would have that a cell only takes input from the neighbours to its left (horizontally and diagonally) and from the neighbour above, we will have a feed-forward motion moving semi-diagonally from the upper left corner to the bottom right. In case the grid does not wrap horizontally and vertically, the effect of a change in state of a cell will never be fed back to it. This means that the simulation is bound to die out (i.e. it will stop changing) after a short period of time, because no new changes will be fed to any of the cells. Although this may produce some patterns that stand out to an observer of the system, this is not what is meant by emergence.

In a slightly different sense, the idea of linked mechanisms is also found in [12], [13] and [14], by Chen et al., in which the authors attempt capture emergence in terms of related events. Their approach is missing the cyclic aspect in their chains of cause and effect, but they do address another important point. They claim that an emergent can never be a simple event, which is the result of application of an elementary rule, but should at least consist of two events that are somehow connected to each other. This corresponds to the idea that something is not emergent if it is explicitly specified in the elementary components of the system, as was already mentioned in section

3.3.

What should further be noted here, is that systems can display many different composite mechanisms, but not all of those mechanisms exhibit emergent features. Whether a macro-level mechanism qualifies to be emergent or not depends on the criteria for emergence that are used (e.g. that the configuration of component states should form a meaningful pattern for the observer). But what we do know is that the cyclic chain reactions captured by these mechanisms lie at the basis of all three categories of emergence.

In conclusion, for a system to produce any form of emergence, it should at least consist of two elementary mechanisms that are capable of forming a macro-level mechanism. In other words, a multi-agent system should have at least two rules A and B where (at least part of) the change resulting from application of rule A, satisfies the condition of rule B (i.e. triggers the application of rule B) which otherwise could not have been applied. Furthermore, it should hold that rule A is part of the rule-set for agent x, rule B is part of the rule-set for agent y, and agents x and y are not the same agent. Rules A and B do not necessarily need to be distinct. Furthermore, the chains formed by rules like these should form a loop. For instance, rule B, that was triggered by rule A, might trigger rule A again. These cyclic chain reactions seem to be one of the most crucial elements for producing emergence.

4.3.2   Initial situation

In [25], Holland briefly mentions that, for a system displaying emergence, changes in initial conditions can have as a consequence that the system no longer displays emergence. The fact that emergence depends on initial conditions, is illustrated by Kubík in [26]. He shows that the Game of Life, known for producing emergent patterns, does not always display emergence. There are configurations of agent states in which none of the agents can apply a rule that results in a state change.

Thus, an important requirement is that the initial situation of a multi-agent system should enable application of rules that bring about some kind of change. To be exact, the initial situation should ensure application of rule A, as specified in 4.3.1, at some point during the simulation.

4.3.3 Choice & Non-determinism

The most simple form of multi-agent system is completely deterministic. By this I mean that for every possible situation, an agent can apply exactly one rule, or none (i.e. he has no choice). This means that every simulation starting with the exact same initial configuration, will give the same results. The Game of Life is an example of such a fully deterministic system, and proves that non-determinism is not needed for producing emergent patterns of category I emergence.

The complexity of a system increases when agents *do* have a choice, in which case an agent needs to be equipped with some kind of decision method. If this decision method incorporates a random element, this can result in variety in behaviour among agents of the same type. On the one hand, this could be at the expense of order and regularity, but on the other hand, it increases the number of different situations that can occur within a system. This implies that non-deterministic elements may increase the chance of category I emergence, or increase the number and variety of instances of category I emergence, in a system.

In case an agent employs a decision method that takes feedback from the global level as (partial) input, the system has potential to produce category II emergence. Note that decision methods like these can be completely deterministic. This suggests that non-determinism might not be necessary for category II emergence either.

4.3.4 Parallelism

In [26], Kubík argues that emergent patterns in the Game of Life depend on the parallelism of the agents' actions. He shows that observed emergent patterns will no longer appear when the parallel activities are replaced by sequential actions. For cyclic patterns consisting of multiple agents, such as the glider, this is rather trivial, since they are *made of* parallel actions. But for other types of phenomena, such as static patterns, this may be different.

If a macro-level phenomenon, produced by multiple agents, can still occur when the agents are replaced by a single agent acting in a sequential manner, Kubík considers that particular macro-level phenomenon to be non-emergent. Even though the behaviour may seem impressive to the external observer of

the system.

Kubík furthermore claims that in systems incorporating stigmergic communication (e.g. the communication through pheromones among ants), parallel actions can be replaced by sequential ones without altering the results. Therefore, he argues that the phenomena produced by these systems do not qualify as emergence.

Although I agree with Kubík that parallel actions should not be replaceable by sequential ones, I do not agree that this is the case with ant colonies. The emergence of shortest paths in colonies of foraging ants *does* depend on parallel actions. Shortest paths emerge because a shorter path can be traversed more times than a longer path in the same amount of time. As a result, the shorter path wil have a stronger trail of pheromones, which will attract more ants. If the actions would be sequential, a trail will not be reinforced enough.

4.4 Organization

By organization I mean the mechanism of the system as a whole that coordinates the temporal order of actions. Many systems follow constant cycles, e.g. each round we first retrieve external inputs, then all agents take action on these inputs, followed by an evaluation of their actions. Chapter 6 will provide a concrete example of a system incorporating this.

4.4.1 Input

The organization of a multi-agent system may include an interface through which a human user or other system can interact with the system at hand. A user could for instance provide input to the system, directly or indirectly influencing the behaviour of agents in the system. If the input is independent of the behaviour of the system (i.e. it is not influenced by the system-behaviour) it adds some perturbation to the system. This could provide the diversity needed for emergence, or it could disturb emergent phenomena such as spatial patterns.

4.4.2 Feedback

The organization of the system is responsible for calculating and providing feedback from a macro level, often the global level (which represents the summed actions of *all* agents), to the local level. This feedback can be seen as a form of indirect interaction between agents and may sometimes even be the only connection between them. It reflects phenomena or outcomes of the joint (parallel) actions of a set of agents. According to section 3.3, feedback plays a role in category II and III emergence.

Chapter 5

Summary so far

The aim of this chapter is to summarize the findings of the first two parts of this thesis, before I continue to the third. In chapter 4 I have identified elements that are important (or even required) for emergence and elements that potentially increase the chance or amount of emergence occurring in a multi-agent system. Findings of the first kind can be used to make the descriptions of the different emergent types formulated in section 3.3 somewhat more concise and detailed at the same time. The one that stood out and seems to lie at the core of emergence, is the concept of a 'chain reaction', as explained in section 4.3.1. This is not very surprising given that it corresponds to the idea of 'feedback', which was already incorporated in the descriptions of the different emergence types in the unified classification. The section on chain reactions in chapter 4 provided a more detailed description of the concept though, and stressed the cyclic nature of feedback. It can furthermore be coupled to another important requirement that can be concluded from chapter 4: parallelism of actions among multiple elementary entities (i.e. agents) of the system. The following is an attempt to capture the most important findings in a single phrase. This could be considered an informal general definition of emergence.

Definition 1. An *emergent phenomenon* in a multi-agent system is a recurrent macro-level phenomenon, noticed by an external observer of the system, that arises through feedback-loops comprising parallel (inter-) actions of (or between) multiple elementary entities in the system.

By recurrent I mean that the phenomenon should occur multiple times

within a single simulation of the system or occur in multiple simulations of the same system. Remember the glider in the Game of Life, that occurs each time a part of the grid displays one of the configurations of the cyclic patterns surrounded by dead cells.

Furthermore, by macro-level phenomenon we mean a phenomenon occurring at some higher level than that of individual elementary entities (i.e. agents), which form the micro-level. It refers to things such as an output produced by, or a configuration of states consisting of, multiple agents. If it involves all agents, I sometimes use the term 'global level' instead of macro-level. Among other things, a phenomenon can be a structure (such as a pattern) or a property (such as a value), possibly characterized by a temporal aspect (e.g. the convergence to a certain value over time).

The concept of feedback-loops can be captured in a more formal definition, which is given below.

Definition 2. A *feedback-loop* involves two mechanisms A and B, where the output of mechanism A serves as (partial) input for mechanism B, and the output of B serves as (partial) input for A. Mechanisms A and B need not be distinct.

The concept of a 'mechanism' as was used here, differs from the concept that was provided by Holland and discussed in section 4.3.1 in that it does not need to be employed by agents. A mechanism can for instance also comprise a calculation of the global outcome of the actions of all agents that determines the global feedback, which is not part of a single agent but is employed by the organizational component of the system. A formal definition of a mechanism is provided below.

Definition 3. A *mechanism M* comprises an array of $n \geq 1$ linked elementary rules of the system, $\{X_1, X_2, ..., X_n\}$, in which the input for mechanism M is taken as input by rule X_1, and the output of the mechanism is the output produced by rule X_n. For every rule X_i in between holds that the output of X_i serves as (partial) input for rule X_{i+1}.

The concepts defined above can furthermore be used to express the three different types of emergence that were treated in section 3.3 in a more formal way than was done in that section.

Definition 4. An emergent phenomenon is considered to be an example of *category I emergence* when it results from a feedback-loop comprising only elementary rules that are employed by multiple individual agents in parallel.

Definition 5. An emergent phenomenon is considered to be an example of *category II emergence* when it results from a feedback-loop comprising a mechanism that is part of the organizational component of the system instead of an individual agent and takes the output of multiple mechanisms as input, producing output (i.e. the global feedback) that is fed to at least one mechanism that is being employed by multiple agents in parallel.

Definition 6. An emergent phenomenon is considered to be an example of *category III emergence* when it results from multiple feedback-loops in which the inner workings of mechanisms (i.e. between input and output) may change over time. Furthermore, a category III emergent phenomenon is capable of directly influencing and interacting with other category III emergent phenomena (i.e. at a macro-level), without having to manipulate the micro-level constituents of the other.

Now I have established a unified classification of emergence and inventoried elements that might play a role in emergence, it is time to test the usability of these findings in practice. This will be done in the chapter that follows.

Chapter 6

Case study: Artificial Financial Market

In this chapter, I will study an implementation of an Artificial Financial Market (AFM) in NetLogo, displaying several phenomena that may be considered emergent. The model used here is based on 'multiagent', a model from the online NetLogo User Community Models library, created in 2006 by Marianna Caldana, Paolo Cova and Umberto Viano [7]. The model simulates a financial market in which a single type of stock is traded. Each round (time step) traders decide whether to buy or sell an asset of this stock. A trader buys a share when he is optimistic ('bullish'), believing the market will rise, and sells one if he is pessimistic ('bearish'). This sentiment of a trader can be influenced by things such as the news or the sentiments of close colleagues. The value of an asset increases when the number of buyers is higher than the number of sellers, and decreases otherwise.

In section 6.1 I will provide a more elaborate description of the model. The changes I have made to the model created by Caldana et al. are documented in the information-tab in the program. These changes have influenced the behaviour of the model, but it still displays phenoma such as bubbles and crashes that are often considered emergent. Since it is my goal to study emergent phenomena rather than simulating a financial market as realistic as possible, I believe the changes I have made are justified. Although I have to note that some changes were driven by my personal intuition on improvement of realism concerning behaviour.

Section 6.2 will contain an analysis of the elements of the AFM-model based on the findings listed in chapter 4. I will indicate which elements of the model could potentially enable emergence.

In section 6.3 I will describe two observed macro-level phenomena that could be regarded as emergent. These phenomena will be related to the different types of emergence as discussed in chapters 2 and 3.

Finally, in section 6.4 I will attempt to connect the macro-level phenomena as described in 6.3 to the micro-level elements of the model. I will examine which of the elements listed in 6.1 played a role in the two selected emergent phenomena.

6.1   Model description

The spatial representation of the market is an $n * m$-grid, where each cell is inhabited by an agent representing a trader. We distinguish three types; an active trader either behaves as a 'fundamentalist', an 'imitator', or a 'monkey'. In the grid, fundamentalists are represented as yellow cells, imitators are blue and monkeys are red. Additionally, black cells represent failed traders, i.e. traders that no longer take part in the market because they have lost too much money. We furthermore indicate whether a trader is optimistic or pessimistic (i.e. whether a trader has just bought or sold an asset) by using two shades of each colour. The lighter shade indicates a pessimistic trader, the darker (brighter) shade an optimistic trader.

The initial distribution of types over the nm traders is user-defined. At the start of every next round, traders can switch between types. A trader of type X considers changing to type Y when the average wealth among neighbours of type Y is higher than that of all other types and $p\%$ higher than his own wealth, where p is user-defined. Whether he will actually change to that type depends on chance, which is a user-defined percentage and equal for every trader that considers to switch types.

After each trader has decided whether or not to change types, the news arrives. The news is an exogenous input, representing information that can possibly influence the market sentiment regarding the stock at hand (e.g. information about the company behind the stock, rumours etc.), taking a value of +1 if positive and -1 if negative. Whether the news is positive or negative

is determined at random.

At this time, traders decide whether to buy or sell an asset. Each type of trader is defined by its distinct decision method. Fundamentalists base their decision on their individual estimate of the value of a share. When a trader becomes a fundamentalist, his estimate equals the current 'log-price' (where $price = e^{log-price}$) of the asset. Each round this trader remains a fundamentalist, a random value between 0 and 1 is added to his estimate when the news is positive, and subtracted otherwise. Fundamentalists decide to buy when they think the value of the share is underestimated, i.e. their estimate is higher than the log-price, and sell otherwise.

Imitators have the most complex decision method. They are influenced by the sentiments of their close colleagues, represented by the eight neighbouring cells in the grid, and by the news. The sentiment of a trader reflects his last trading decision. If he bought a share, his sentiment-value is 1, if he sold a share, it is -1. The sum of the sentiments of all neighbours of type X is divided by the number of neighbours of type X, which is called the X-return. For each trader type $X \in \{Imitator, Fundamentalist, Monkey\}$, the X-return is weighted by an imitator's sensitivity to the sentiment of traders of type X, and the news is weighted by an imitators' sensitivity to the news. Sensitivity-values always range between 0 and 1, the news is either 1 or -1, and the returns range between -1 and 1. If the sum of the weighted returns of each trader type and the weighted news is positive, an imitator will decide to buy. If it is negative, it will sell. And if it is zero, the imitator will decide at random.

The third trader type, the so-called monkeys, are the only traders that are not influenced at all. They decide at random whether to buy or sell a share.

Initially, each trader starts with a portfolio holding one share and a wallet containing a user-defined amount of money (i.e. endowment, which is later called liquidity). The performance (or wealth) of a trader is measured by his balance, which is the sum of his 'liquidity' and his 'portfolio-value'. The first is the amount of money a trader has at his disposal. The latter is the total value of the shares currently owned by the trader, calculated by multiplying the number of shares with the current price of a share.

When all traders have made their decisions, their liquidity and portfolio-value is updated, as well as the price of the asset. The price of the asset

is increased by the return. The return of a set of agents A, is the number of agents in A that were buyers minus the number of agents in A that were sellers, divided by the number of active agents in A. Thus, the return is positive if the number of buyers is higher than the number of sellers, and negative otherwise.

Then, the balance of each active trader is evaluated. If the balance of a trader is below a user-defined maximum debt, the trader is considered failed and can no longer take part in the market. If his liquidity is below the maximum debt, but his balance is not, a trader is forced to sell shares until his liquidity is acceptable, but only as long as his balance remains above the maximum debt. If a trader didn't have enough assets to pay enough of his debt, he will be considered failed. A third possible scenario is when a trader has a shortage of assets, i.e. his balance is above the maximum debt, but his portfolio-value is not. In this case, a trader is forced to buy shares until this shortage is resolved, unless he does not have enough money to accomplish this, in which case he is considered failed.

Finally, the sensitivity-values of each imitator are updated. When a trader becomes an imitator, he is initialized with a random sensitivity-base between 0 and a user-defined maximum, smaller than or equal to 1, for each type of trader as well as for the news. At the end of each round, an imitator updates his sentiment-sensitivity to traders of type X by taking the sensitivity-base and adding or subtracting the return of his neighbours of type X. An imitator becomes more sensitive to the sentiments of traders of type X when the return of his neighbours of type X reflects the market-movement, i.e. if both the return of the market as a whole and the return of the neighbours of type X is positive (or negative). At this point, the news-sensitivity of an imitator is updated as well. The news-sensitivity ranges between zero and a user-defined maximum of at most one. If the movement of the market corresponded to the value of the news, his news-sensitivity is increased by the global return, otherwise it is decreased by the same value. Thus, if positive news was followed by a positive return, an imitator will become more sensitive to the news.

6.2 Elements potentially enabling emergence

In chapter 4, I have listed several elements that could possibly enable emergence in multi-agent systems. In this section I will examine which of these elements are incorporated in the AFM-model. This section globally has the same structure as chapter 4, except that I have merged a few topics here.

6.2.1 Environment

The environment in the AFM-model plays no active role in the system. It only provides a spatial structure to facilitate connections between agents. This structure is a static two dimensional grid with periodic boundary conditions, of which the dimensions can be defined by the user. Each cell in the grid represents an agent, and each cell is linked to its eight surrounding neighbours. Recall that in the Game of Life, this similar setting played a role in the coming about of category I emergence. I will get back on this in the next section when I discuss the connectivity between agents.

6.2.2 Connectivity & dependence

Agents in the AFM-model are immobile and permanently connected to their eight surrounding neighbours. Each round, a trader has to decide whether or not he switches types, and if so, to which type. For this decision, he is dependent on information concerning his active neighbours. Thus, this decision is based on local feedback only. On top of this, imitators are also dependent on their active neighbours in their decision whether or not to buy an asset. I will elaborate on this in section 6.2.7.

The fact that traders can change types based on local feedback could lead to patterns in the spatial distribution of different types of traders over the general population that are noticed by an external observer of the system. Hence, this is a form of category I emergence that could be produced by the system.

Agents that become inactive no longer take part in the market; neighbouring agents no longer include information about this agent in their decision making processes. Thus, when agents become inactive, connectivity decreases. And connections that are lost will never be reinstated since a trader that has failed is *permanently* excluded from the market.

When the number of failed traders in the neighbourhood of a trader increases, the extent to which the neighbourhood reflects the overall population will decrease. Like in statistics, the smaller the sample size, the less accurate the estimate of the population as a whole. As a consequence, there might be an increase in type-changes. A trader could for example switch to a type that is represented by only one of his neighbours, who could have become wealthy 'by accident', because this peak is not smoothed out. On the other hand, when many traders have failed, the remaining active traders might have less diversity in trader types among their neighbours which implies less type-changes. All in all, the decrease in connectivity due to failing traders will have some effect on the behaviour of the remaining active traders. This may influence the patterns considered to be of emergence category I, or result in new emergent patterns.

6.2.3 Quantity

Financial markets typically consist of a large population of traders. In our model, the market is represented by a 51x45-grid, initially containing 2295 active traders. The dimensions of the grid can be adjusted by the user, but should not become too small, since I argued in section 4.2.2 that this would decrease the possibilities for emergence. Over time, the number of traders participating in the market may decrease when traders fail and get excluded from the market.

We can describe three dynamic sub-populations among traders, based on the three different types of traders. (See the corresponding plot 'Distribution of types' in the interface of the model.) The change in size of each population over time, due to individual decisions based on local feedback, may display certain patterns. Although these changes affect local behaviour, they do not play an explicit role in the system, and would therefore only qualify for category I emergence.

6.2.4 Heterogeneity

In section 4.2.3, I showed that heterogeneity is not necessary for emergence, but I argued that it could increase the chance of emergence. In the AFM-model, the user determines whether the initial population of agents regarding

trader types is heterogeneous. If it is homogeneous, it will remain homoge-neous, because an agent can only switch to a type that is represented by at least one of his neighbours. In this case, it is highly unlikely that we will observe any category I emergence in the from of spatial patterns. Monkeys decide at random whether to buy or sell, so we will not observe any regularity in the spatial patterns of buyers (deep red) and sellers (lighter shade of red). Fundamentalists are also independent of their neighbours and will therefore not produce regular patterns either. Imitators will display homogeneous be-haviour provided that all agents have a non-zero news-sensitivity value. In the first round they will all rely fully on the news. Sentiments play no role in the first round since traders are initialized with neutral sentiments. This ensures that the movement of the market corresponds to the decisions of the imitators, which increases the sensitivity to the sentiments of other imitators by 1, reinforcing homogeneous behaviour.

Since this these behaviours can be predicted by simply inspecting the micro-level elements of the system, these behaviours do not qualify for emer-gence. In case the initial population of traders is heterogeneous regarding trader types it is more difficult to make predictions because the number of possible interactions is much higher than for homogeneous populations. Thus, a heterogeneous initial population opens up a possibility for emergence to occur in the system.

Traders of the same type are homogeneous with respect to their decision methods, i.e. each agent of the same type employs the same decision method. However, the resulting behaviour can be heterogeneous due to the fact that input values for these decision methods differ among traders. Fundamentalists keep an individual estimate of the value of an asset, on which they base their decision whether or not to buy. The decision method of imitators makes use of individual sensitivity-values with respect to the news and to the sentiments of others. And finally, monkeys decide at random whether they will buy or sell a share, which is based on a random number between 0 and 1 that is different for each monkey. This heterogeneity among traders of the same type increases the chance of emergence.

6.2.5 Decentralization

In the AFM-model, both knowledge and control are decentralized. There is no agent or entity that tells others what to do and agents only have access to information about their neighbours and partial knowledge of the global level (i.e. the global return and price). In 4.2.4 I argued that this is necessary for any type of emergence. Thus, this requirement is met by our model.

6.2.6 Adaptivity

In section 4.2.5 was indicated that the presence of adaptive agents opens up the possibility of category II emergence ocurring in a system. Imitators in the AFM-model are an example of such adaptive agents. Their behaviour is influenced by feedback from the global level in combination with local feedback. If the news corresponds to the movement of the market, an imitator will become more sensitive to the news. This happens when the news that arrived at the start of a round was positive (negative) and was followed by an increase (decrease) of the price of an asset. Otherwise, he will become less sensitive to it. The same holds for an imitators' sensitivity to the sentiments of the different trader types. If the return of his neighbours of type X corresponds to the movement of the market (i.e. the global return), an imitator becomes more sensitive to the sentiment of traders of type X. Otherwise, he will become less sensitive to their sentiment.

6.2.7 Chain reactions

In section 4.3.1 I emphasized that chain reactions are important for emergence. In the AFM-model, the design of the system forces each active trader to take action every round. The change resulting from the application of a rule by agent x, does not trigger application of a rule in neighbouring agent y, but it *can* serve as an input for the rule (decision method) agent y is forced to execute by the system. So in some sense, an agent reacts to his neighbours, creating a slightly different kind of chain reaction than was treated before.

Three major (cyclic) chain reactions can be distinguished in the AFM-model. The first chain reaction results from the fact that, at the start of each round, an active trader is forced to choose whether or not he will change types. Each trader uses the wealth of his active neighbours as partial input

for the universal method for making this decision. The wealth of each trader is the result of his cumulated actions and thus can be seen as an output of his actions. The decision of a trader concerning his type determines the decision method concerning trading for that round, which will have an effect on his wealth. And his wealth is taken as a partial input by his neighbours in the next round. The resulting chain is depicted in figure 6.1. These actions take place at a local level (i.e. the method only retrieves information from neighbours), which points towards a possibility for category I emergence.
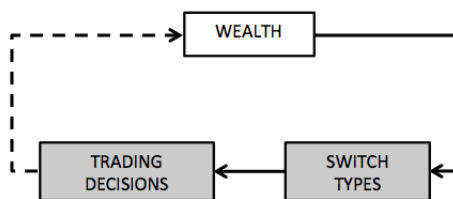
Figure 6.1: Diagram depicting the chain resulting from the switch-types procedure.

In the diagram above, and the ones that will follow, coloured boxes represent methods employed by traders and white boxes are outputs of those methods or of system-methods. The latter are methods such as the procedure to calculate the price from the return. These system-methods themselves are not depicted.
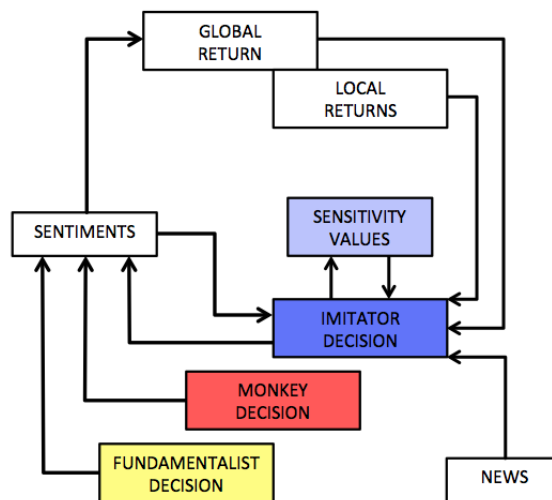
Figure 6.2: Diagram depicting the chain involving imitators.

The second chain reaction arises through the decision method of imitators

regarding selling or buying an asset. An imitator takes the output (i.e. the sentiments) of the decision methods employed by his neighbours in the previous round as (partial) input for his decision method in the current round. The output of his decision is used as an input in the same way by the imitators among his neighbours in the next round. This, again, is a mechanism that takes place at a local level, and therefore points towards category I emergence.

However, there is another element that plays a role in this second chain reaction. As already mentioned at the end of the previous section, the decision method of imitators also includes an update method for sensitivity-values. This update method partially depends on feedback from the global level (i.e. the movement of the market as a whole), which is an indicator for category II emergence. The complete chain involving imitators is depicted in figure 6.2.
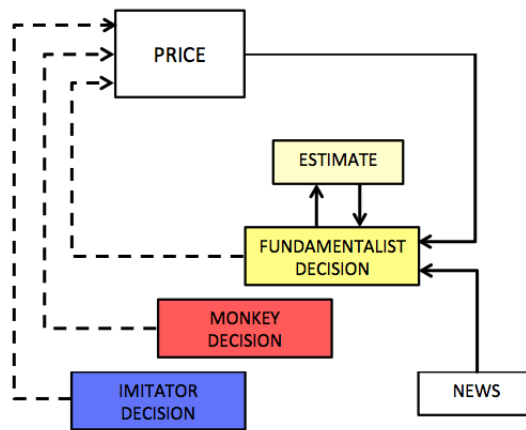
Figure 6.3: Diagram depicting the chain involving fundamentalists.

The third chain reaction involves fundamentalists and is depicted in figure 6.3. In their trading-decision, fundamentalists compare the current price to their individual estimate. The current price is the result of the combined decisions of all traders in the previous round. Thus, the decision method of a fundamentalist takes global feedback as partial input. The other part of the input is the news, which was an exogeneous input. The news is used to update the personal estimate of the price of a share, before it is compared to the actual current price. The resulting decision of the fundamentalist influences the price of the share in the current round, which is taken as partial input by fundamentalists in the next round. Like the chain involving imitators, this chain involving fundamentalists might lead to category II emergence. It could

for instance manifest itself in the form of recurrent patterns or tendencies in the course of the price.

The three cyclic chains partially overlap, meaning they are dependent on each other. The three chains are merged in figure 6.4, providing a complete picture of the chains in the system. The fact that multiple traders act in parallel adds to the complexity of the structure of the system as depicted in this diagram. It is clear that even when we would have only a handful of traders in the system it would already be hard to predict what will happen exactly.
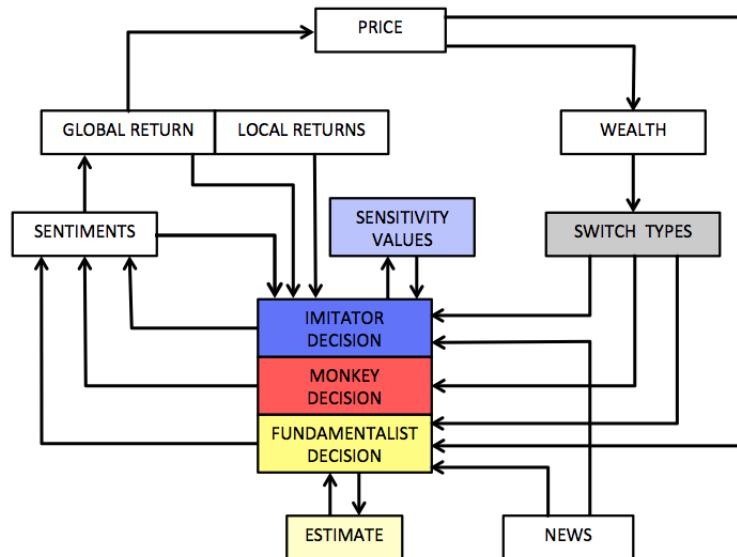
Figure 6.4: Diagram depicting all cyclic chain reactions in the model.

6.2.8   Initial situation

In section 4.3.2 I showed that the initial situation might determine whether or not emergence will occur in a system. In the AFM-model, every trader is forced to make a decision each round, and the resulting decisions change the price of a share. Thus, given any initial situation, rules will be applied that bring about some kind of change. These changes are then used as (partial) input for other rules. This implies that the initial situation will never block emergence. It can however constrain the *forms* of emergence that may occur. For instance, emergent phenomena in which three types of traders are involved cannot occur when only two types are present in the initial situation.

6.2.9 Parallelism

In section 4.3.4 I claimed that parallelism is one of the requirements for emergence. NetLogo can not handle multiple actions at the same time; if a set of agents has to execute a procedure, they will always do so one by one. But we *can* simulate parallelism. We just have to make sure that decisions regarding trading are based on the situation at the end of the previous round only, and not partially on decisions already made in the current round. To accomplish this, we let all agents decide before the first agent actually changes his sentiment. Their decisions are indicated by a change in the number of shares in their portfolio. As long as the portfolio-values are updated after each trader has made a decision, this has no influence on the decisions of other traders.

The resulting decisions would naturally be different when actions would be entirely sequential, i.e. when sentiments are updated at the moment of decision. Suppose you are an imitator and it is your turn to decide whether to buy or sell a share. If one of your neighbours has preceded you, you will retrieve information about the sentiment of this neighbour (as input for your decision method) that is already one round ahead of you (and probably some of your other neighbours). This might have as a consequence that emergent phenomena occurring in the parallel situation will not occur when actions are sequential.

6.2.10 Non-determinism

In section 4.3.3 I argued that non-determinism is not necessary for emergence, but the variety resulting from non-deterministic elements might increase the chance of emergence. The AFM-model contains several non-deterministic elements ensuring variety among traders of the same type. One that might matter, is the increment- or decrement value of fundamentalists updating their personal estimate of the price of a share, which ranges between 0 and 1. In the original AFM-model by Caldana et al., the estimate was a global variable, which resulted in homogeneous behaviour among fundamentalists. This meant that when the majority of traders was a fundamentalist, they would determine whether the price of the share would increase or decrease. So this affects the course of the price, and also the wealth of traders. In some sense, it would constrain some of the possible emergents in the system.

Another non-deterministic element is the set of sensitivity-values of an imitator regarding other traders' sentiments and the news. The user defines the maximum news-sensitivity of an imitator, which is a value between 0 and 1. Each trader becoming an imitator will be initialized with a random news-sensitivity ranging between zero and this maximum. The value of the news is weighted by an imitators' sensitivity to the news in his decision method concerning trading. Another input to this decision method is, for each of the three types of traders, the return of neighbours of type X, weighted by a traders' sensitivity to the sentiment of type X. This sensitivity is initially set to a random value between 0 and a user-defined maximum (ranging between 0 and 1). If the sensitivity-values for each imitator would be the same, the decisions of imitators will not necessarily be homogeneous because each trader has different neighbours.

Furthermore, at the end of each round, the sensitivity-values are updated, representing a valuation of the behaviour of neighbours per trader type and whether or not the movement of the market followed the news. The sensitivity-values regarding the sentiments of others are reset to the initial value and increased or decreased by the return of the neighbours per trader type. Thus, the sensitivity-values differ among imitators. If they would be the same for every imitator, it would still depend on the sentiments of the neighbours whether the imitator will decide to buy or sell, so there could still be variety in decisions of imitators.

6.2.11   Input

The news is an 'exogenous' input to the system. Although it is generated within the system itself, its value does not depend on any process within the system but represents fictitious information from outside the system. The value of the news is used in the decision procedures of imitators and fundamentalists. The news is a random value but it is the same for each trader. This may lead to behaviour that seems random but at the same time homogeneous in some sense.

6.3 Observed emergent phenomena

In the previous section I have inventoried which forms of emergence we could potentially observe in the AFM-model based on the elementary ingredients of the system. I will now describe two phenomena that are actually observed in the model and can be considered emergent. Based on section 3.3 I will furthermore indicate which types of emergence are possibly represented by these phenomena.

6.3.1 Maze-like patterns

The first phenomenon selected for study within this chapter occurs in the spatial representation containing different colours to represent different trader types and their sentiments. In simulations of the AFM-model we can frequently observe maze-like structures, emerging from a completely random initial configuration of trader types. Figure 6.5 shows states ocurring after 50 rounds for simulations with different initial ratios of trader types. It suggests that the maze-like patterns emerge regardless of the initial distribution of trader types, although they are more explicit when one trader type is represented less than the other two in the initial configuration.

The maze-like structures are global patterns that attract the attention of an external observer because they display a certain regularity and are recurrent. Naturally, the larger the population of traders, the sooner an observer will notice the regularity. The patterns will probably not be noticed when we zoom in to a piece of 3x3 grid. This means that we have to observe the model at a suitable macro-level to notice the patterns.

The emergence of these maze-like patterns coincides with a decrease in the number of traders switching types. This results in a phase in which the global pattern seems to be rather stable. However, it is never completely stable. Every once in a while clusters of changing traders move over the grid. And at some point, the patterns may be disturbed and become less regular. This decrease in regularity often coincides with an increase of failed traders. But this does not always occur. There are simulations of the model with only imitators and monkeys in which the maze-like patterns will never be lost.

To the observer there appears to be some organization among traders regarding their types, that is not obvious given the elementary rules of the sys-

tem. It is not explicitly stated anywhere in the code that traders should form these patterns, thus it must result from some kind of chain reaction composed of elementary rules.
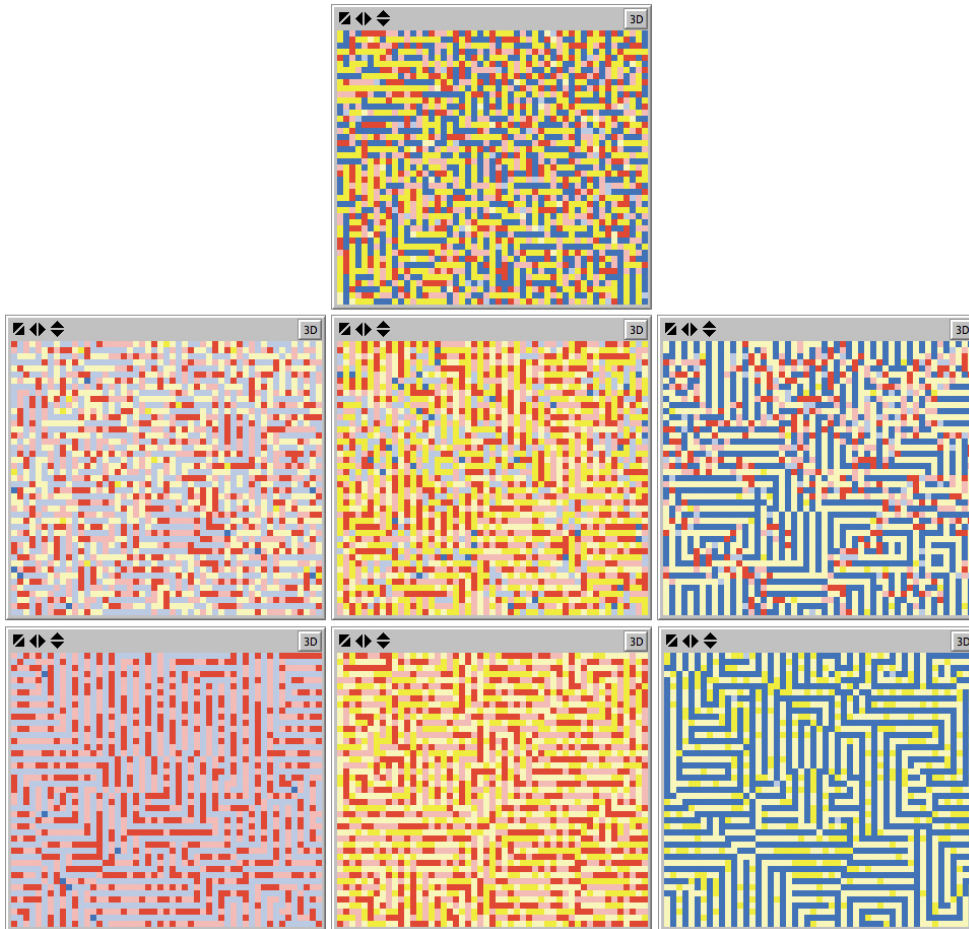
Figure 6.5: States of simulations at t = 50 with different initial ratios of trader types. Initial ratios of trader types from left to right (fundamentalists : imitators : monkeys) are: $1 = (\frac{1}{3}{:}\frac{1}{3}{:}\frac{1}{3})$, $2 = (\frac{2}{10}{:}\frac{4}{10}{:}\frac{4}{10})$, $3 = (\frac{4}{10}{:}\frac{2}{10}{:}\frac{4}{10})$, $4 = (\frac{4}{10}{:}\frac{4}{10}{:}\frac{2}{10})$, $5 = (0{:}\frac{1}{2}{:}\frac{1}{2})$, $6 = (\frac{1}{2}{:}0{:}\frac{1}{2})$, $7 = (\frac{1}{2}{:}\frac{1}{2}{:}0)$. (The remaining parameters were held constant; change-threshold = 0.33, change-probability = 0.7, maximum-debt = 0, endowment = 270, all max-sensitivity-values = 1.0.)

In section 6.2.7 I already predicted that the chain reaction invoked by the switch-types procedure might produce emergence in the form of spatial patterns. The switch-types method uses only local feedback as input. This local feedback contains values that result from other processes involving feedback from the global level. In other words, the switch-types chain as was depicted

in figure 6.1 is part of a larger chain. However, the switch-types method itself, that is considered to be mainly responsible for the emergent patterns, does not use global feedback as an input.

Besides exhibiting regularity, the patterns have no functional meaning to an external observer. They do not provide some solution, and they do not yield useful insights into the workings of a financial market. And, although the spatial configuration of trader types may influence the global behaviour of the traders that determine the price of the asset, the emergent patterns do not have an explicit function in the system either.

In conclusion, it seems that the maze-like patterns meet the criteria specified in definition 4 as provided in chapter 5, and thus can be considered to be a form of category I emergence.

6.3.2 Bubbles & crashes

The second potentially emergent phenomenon selected for study is the occurrence of sudden changes of tendencies in the course of the price of a share, known as bubbles and crashes. Figure 6.6 provides an example. We see for instance that the global trend of the price is increasing over rounds 200 to 350 (i.e. a bubble), but suddenly decreases after that (i.e. a crash), until it starts increasing again around round 500. This behaviour cannot be reduced to any of the elementary rules of the system, i.e. there is no elementary rule responsible for creating bubbles and crashes.

The reason for incorporating this phenomenon is that it might be an example of category II emergence. The price of a share results from the combined decisions of all traders in the market, and some of these decisions are based on the previous price of the share. The trend in the price of a share may thus influence the local behaviour. This reveals a feedback-loop involving global feedback, which is typical for category II emergence.

Bubbles and crashes in a market are generally claimed to be emergent by those who observe them. The reason why people often consider these phenomena to be emergent, is that they are not coordinated by some leader, but that they are the result of the combined individual decisions of all participants in the market. These phenomena have already been extensively studied within what is called 'behavioural finance theory', which yielded many theoretical explanations. See for instance [16]. However, I am not interested in social causes

(such as herding behaviour) and meaning here. In this chapter I will focus only on the pure underlying mechanisms that are specific for our implementation.
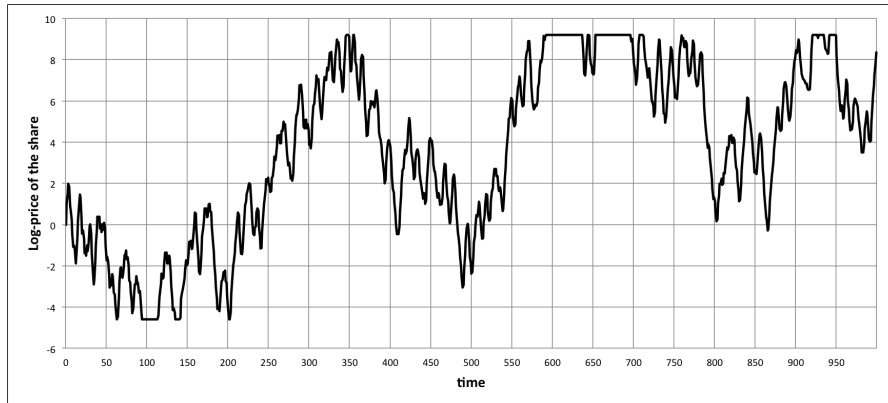
Figure 6.6: Example of the sudden shifts in the tendency of the price of a share. Initial settings of the model: 50% Imitators, 50% Fundamentalists, change-threshold = 0.10, change-probability = 0.70, maximum-debt = 0, endowment = 300, all max-sensitivity-values = 1.0.

6.4 Analysis of emergent phenomena

In section 6.2 I discussed which elements of the AFM-model could potentially lead to, or increase the chance of, emergence in this model. In this section I will attempt to reveal which of these elements are responsible for the specific phenomena selected for study in this chapter and analyze how the observed macro-level phenomena resulted from the elementary elements of the system.

6.4.1 Maze-like patterns

In section 6.3.1 I have already mentioned that the switch-types procedure is the main part of the mechanism underlying the emergence of the maze-like patterns. Remember the loop from section 6.2.7 involving the switch-types procedure, of which a copy is depicted in figure 6.7 below. If one trader switches types, it could cause one of his neighbours to change in the next round. The following example will illustrate this.

Consider the situation in table 6.1[1]. The start of round t represents the situation in round t before any trader has switched types, i.e. the situation

---

[1]Colours in the table do not refer to specific trader types of the AFM-model.
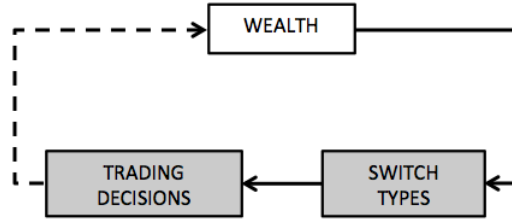
Figure 6.7: Diagram depicting the chain resulting from the switch-types procedure.

at the end of the previous round (round $t - 1$). If trader 4 would change types in round t, this would result in the situation depicted in table 6.2. This represents the situation after all traders had the chance to switch types in round t but before any trader has switched types in round $t + 1$. Now let the wealth of trader i at the start of round t be denoted by $w_i(t)$. Furthermore, let the total wealth of yellow neighbours for trader i at the start of round t be $Y_i(t)$, and the wealth of blue neighbours be $B_i(t)$.

Table 6.1: Start of round t

Table 6.2: Start of round $t + 1$

The wealth of trader 4 at the start of round t is 5, whereas the wealth of each other trader is 10. Let us assume that all 8 traders buy a share in round t, increasing (or decreasing) their wealth by the number of shares they had at the start of round t multiplied by the increase (or decrease) of the price at the end of the round. Furthermore assume that the price of the share increased by a value q, each agent has an equal amount of money, and that the portfolio of trader 4 holds 5 shares while all other traders each have 10 shares in their portfolio at the start of round t. Then we have that:

$$B_5(t) = w_1(t) + w_3(t) + w_4(t) + w_7(t) + w_9(t) = 45$$
$$Y_5(t) = w_2(t) + w_6(t) + w_8(t) = 30$$

$$\overline{B}_5(t) = \frac{B_5(t)}{5} = 9$$

$$\overline{Y}_5(t) = \frac{Y_5(t)}{3} = 10$$

$$\overline{B}_5(t+1) = \frac{B_5(t) + 40q - w_4(t)}{4} = \frac{40 + 40q}{4}$$

$$\overline{Y}_5(t+1) = \frac{Y_5(t) + 35q + w_4(t)}{4} = \frac{35 + 35q}{4}$$

This means that $\overline{B}_5(t) < \overline{Y}_5(t)$, while $\overline{B}_5(t+1) > \overline{Y}_5(t+1)$. Thus, trader 5 will consider switching types in round $t+1$ due to the change of trader 4. If he will do so, this could trigger a type-change in any of *his* neighbours in the next round. So this is a clear example of a chain reaction as was meant in section 4.3.1.

Recall that in section 6.2.7 I have furthermore indicated that the switch-types chain reaction is part of a bigger mechanism, since the method for switching types uses the wealth of neighbouring traders as input. The wealth of each trader is dependent on the price of a share (and the amount of shares he has in his portfolio), which depends on the global return. The global return is determined by the summed trading-decisions of all active traders. And I have shown in 6.2.7 that imitators and fundamentalists had their own loops. The trading-decisions of fundamentalists depend on the price of the share and the decision of imitators partially depend on the returns. The complete loop is repeated in figure 6.8.

I have now identified the chains that lie at the base of the emergent patterns. However, this does not yet provide an explanation for the shapes of the patterns. This requires a mathematical analysis of these chains, based on a detailed examination of the underlying code.

On close inspection of the code defining the switch-type procedure, it turned out that it contains an inaccuracy in the gathering of input for the switch-type decision method. By mistake, the wealth of a trader type among the neighbours of a trader is added to the wealth of that type of the previous round. Thus, this wealth is accumulated. To obtain the average wealth per trader type among neighbours, the wealth is divided by the current amount of traders of that type among neighbours. So the 'average' wealth calculated in the code is not the actual average wealth of a trader type among a traders' neighbours in the *previous round* only. As a consequence, the wealth of a type
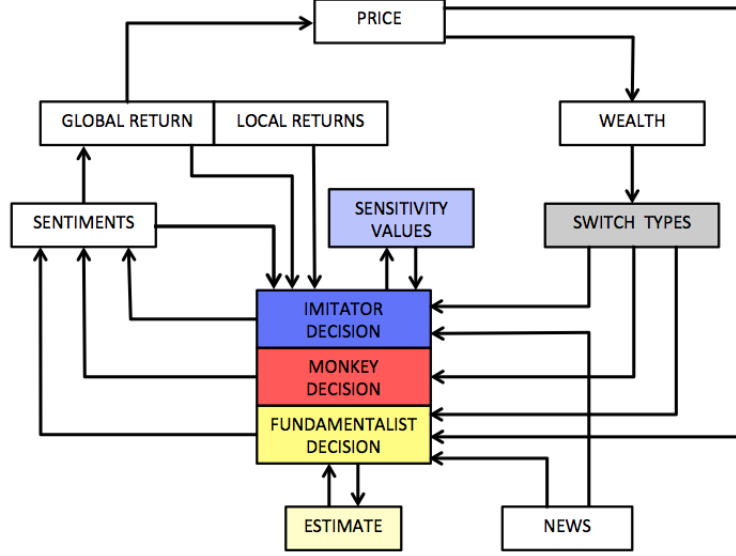
Figure 6.8: Diagram depicting all cyclic chain reactions in the model.

will always increase provided that that type is represented by at least one neighbour, while the number of traders of that type always remains between 0 and 8 (inclusive).

The formal notation of this calculation is given below, where $N_x(t)$ is the set of traders of type x among the neighbours of a trader at the start of round t, $n_x(t)$ is the number of traders of type x among neighbours (i.e. the cardinality of $N_x(t)$), $w_i(t)$ is the wealth of trader i at the start of round t, $sw_x(t)$ is the summed wealth of traders of type x among the neighbours of a trader at the start of round t, $cw_x(t)$ is the cumulated wealth of type x among neighbours (thus over all previous rounds including the wealth at the start of round t), $aw_x(t)$ represents the erroneously calculated 'average' wealth, and $\overline{w}_x(t)$ is the (correct) average wealth of type x among neighbours at the start of round t.

$$sw_x(t) = \sum_{i \in N_x(t)} w_i(t)$$
$$\overline{w}_x(t) = \frac{sw_x(t)}{n_x(t)}$$

$$cw_x(t) = cw_x(t-1) + sw_x(t)$$

$$aw_x(t) = \frac{cw_x(t)}{n_x(t)}$$

$$= \frac{cw_x(t-1)}{n_x(t)} + \frac{sw_x(t)}{n_x(t)}$$

$$= \frac{cw_x(t-1)}{n_x(t)} + \overline{w}_x(t)$$

Division by zero will return zero, thus the only case in which the average wealth of a trader type is correct, is when there were no neighbours of that type in the previous round.

We can correct this error in one of two ways. We could eliminate the cumulation of wealth, or leave it unaltered and let the amount of neighbours per type accumulate as well. In the first case, a trader will base his decision on information of the previous round only, whereas in the second case he will take into account the achievements of a type over all past rounds. For the first solution we replace $aw_x(t)$ by $\overline{w}_x(t)$. For the second, we introduce $cn_x(t)$, as defined by equation 6.1, representing the cumulated number of traders of type x among neighbours over all previous rounds including t. Then we replace $aw_x(t)$ by $\overline{cw}_x(t)$, which is defined by equation 6.2 below. Both solutions have as a consequence that we will no longer observe the maze-like patterns.

$$cn_x(t) = cn_x(t-1) + n_x(t) \tag{6.1}$$

$$\overline{cw}_x(t) = \frac{cw_x(t-1)}{cn_x(t-1)} + \overline{w}_x(t) \tag{6.2}$$

Although the 'error' in the code undermines the realism of the maket model, it does not necessarily imply that the maze-like patterns can no longer be considered emergent. As I have already mentioned, the aim of this experiment is not to reveal social causes or discuss the meaning of the emergents, but to focus exclusively on the underlying mechanisms. Indeed, the patterns are now easier to explain, which I will show shortly, but they still meet the requirements for category I emergence as defined in chapter 5 (see definition 4). The patterns still result from chain reactions and are not explicitly specified by elementary rules.

The wealth per type among neighbours will increase as long as it is represented by at least one of a traders' neighbours. After a certain amount of rounds, the erroneously calculated wealth of each neighbouring trader type

will be ridiculously high, to the extent that even division by 8 (all neighbours are of the same type) yields an average wealth that is still higher than the traders' own wealth. Suppose we are at some point in time where this is the case. Let p be the probability that a trader will switch types (i.e. the change-probability) when this seems beneficial to him. Then a trader will always switch types with probability p, given that the neighbouring trader type with the highest 'average' wealth is not the same type as his own current type. At this time, it is more likely that a trader will switch to the type that is currently represented the least among his neighbours, since this type will have the highest 'average' wealth.

This explains the fact that in the maze-like patterns, traders of type X often have two neighbours of the same type (i.e. $n_X = 2$) and six neighbours of type Y (i.e. $n_Y = 6$). However, this does *not* yet explain why the two neighbours of the same type most of the time are situated to the traders' left and right, or above and below the trader, forming a line.

I will now consider several possible situations, given that we are at some point in time where the (erroneously calculated) 'average' of each neighbouring trader type will always be higher than an individual traders' wealth. For simplicity, I assume that the simulation at hand only contains two trader types.

1	2	3
4	5	6
7	8	9

Table 6.3: Situation 1

1	2	3
4	5	6
7	8	9

Table 6.4: Situation 2

Take a look at the situations illustrated in tables 6.3 and 6.4, and assume that the types of all traders surrounding these 9 depicted traders are the same for both situations. In situation 2, the tendency to switch types for traders 4,6,7 and 8 is lower than in situation 1 because their number of neighbours of the similar type is lower in situation 2, which means that the 'average' wealth among neighbours of the same type will be higher than in situation 1. Only one trader, trader 3, has a higher tendency to switch types in situation 2 than it has in situation 1. For the other traders (i.e. 1, 2, 5 and 9), the ratio between neighbour-types remains the same. This implies that situation 2 is

more stable than situation 1.

We can assume that the chance of a trader considering to switch types when his own type is represented by two neighbours while the opposite type is represented by six neighbours, is negligible. Consider trader 5 in tables 6.3 and 6.4. In both situations, trader 5 has two neighbours of the similar type, and six neighbours of the opposite type. This means that he will only consider to switch types when the cumulated wealth of the opposite type is at least three times higher than the cumulated wealth of his own type. Since I assumed that we are at a time where values for cumulated wealth are extremely high, this means that there should be a huge difference between the cumulated wealth the two types, which is highly unlikely.

For a trader with *one* neighbour of the similar type, and seven of the opposite type, the difference needs to be even higher (i.e. at least $3\frac{1}{2}$ times higher). So it is quite unlikely that he will change types. For a trader that has *three* neighbours of the similar type, and five of the opposite type, the cumulated wealth of the opposite type has to be at least $1\frac{2}{3}$ times higher than the cumulated wealth of neighbours of his own type. This difference is less significant, hence the chance that a trader in this situation will consider to change types is reasonable. For a trader that has four neighbours of the similar type, and four of the opposite type, the cumulated wealth of the opposite type simply has to be higher than the cumulated wealth of neighbours of his own type. Traders with more than four neighbours of the similar type will always consider to switch types.

Table 6.5: Situation 3

Now take a look at table 6.5, in which trader 5 only has one neighbour of his own type and seven of the opposite type. This trader has the lowest possible non-zero probability that he will switch types. However, this means that neighbouring trader number 4, which is of the opposite type, has at least four neighbours of *his* own type. For traders 2 and 8, this is at least three. Obviously, this is an unstable situation. It is most likely that trader 4 will

switch types, creating the more stable situation that was already depicted in table 6.4. In the 'worst' case, the four neighbours of trader 4 that were not depicted in table 6.5 are yellow. Then, if the cumulated wealth of blue traders is higher than that of the yellow ones, it is certain that trader 4 will not switch types. If the cumulated wealth of the yellow ones are higher, there is a $p\%$ chance that trader 4 will switch types.

A configuration that would result in an equilibrium (given that we are far enough along in our simulation) is depicted in table 6.6. In this situation, every trader has two neighbours of the similar type, and six neighbours of the opposite type. Because of periodic boundary conditions, the grid should contain an odd number of rows. Now as long as differences in cumulated wealth among neighbours are small (i.e. below the point that the cumulated wealth of traders of the similar type is three times lower than that of the opposite type) for each trader, this configuration will persist. However, it is highly unlikely that this configuration will arise in the first place. All traders have to change to this configuration at once, because if there is at least one trader that did not follow, this trader will disturb the others around him. Since we start from a random distribution, the chance that this will happen is almost nonexistent. And when we would start from an ideal configuration, it will immediately be destroyed since the cumulated wealth of neighbours is not high enough yet.
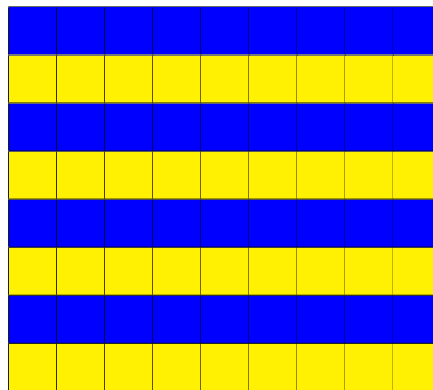

Table 6.6: Equilibrium state

Even when this situation *would* occur at some point during the simulation, it will not last indefinitely. When a yellow trader has two yellow neighbours and six blue neighbours, the cumulated wealth of the blue neighbours will most likely rise faster than that of the yellow neighbours. If this situation

persists long enough, the cumulated wealth of blue neighbours will eventually reach a point at which it is over three times as high as the cumulated wealth of the yellow neighbours, which was the point at which a trader will consider to switch types again.

So now I have revealed the main causes of the maze-like patterns. In hindsight, after discovering the slipup in the code, this turned out to be not that complicated. If I had spotted the issue concerning the updating of the wealth before running the first simulation, I could have calculated the consequences of the error in the same way I did now when explaining the patterns. So in theory, I would have been able to predict the maze-like patterns beforehand, if I had known where to look.

For some this might be a reason for considering the patterns as non-emergent. However, this is undesirable, since it implies that a phenomenon can be considered emergent at first, while the same phenomenon can be considered non-emergent later. Or it can be considered emergent by one observer and non-emergent by another observer at the same time. The concept of emergence adopted in this thesis is dissociated from this issue of subjectivity. The patterns, although explained, still meet the criteria specified in the definitions constructed in chapter 5, thus remain to be considered emergent.

6.4.2 Bubbles & crashes

The analysis of the bubbles and crashes will be kept brief, since the analysis of the maze-like patterns was rather elaborate. The feedback-loops that might play a role in the emerging trends in the course of the price are the chain reaction depicted in figure 6.9 involving fundamentalists, and in a slightly different sense the chain reaction captured by figure 6.10 involving imitators. (The figures were copied from section 6.2.7.) The latter does not directly include the price of a share, but it does include the global return which directly correlates to the price. In the model that was used to construct this one, Caldana et al. already observed that the crashes and bubbles occurring in their model are "(...) normally due to an unusual dominance of one mood over the other in the imitators".

At some point during this analysis, a plot containing the cumulative news-value was added to the model. This revealed a strong correlation between the value of the price and the cumulative news-value. This is illustrated by figure
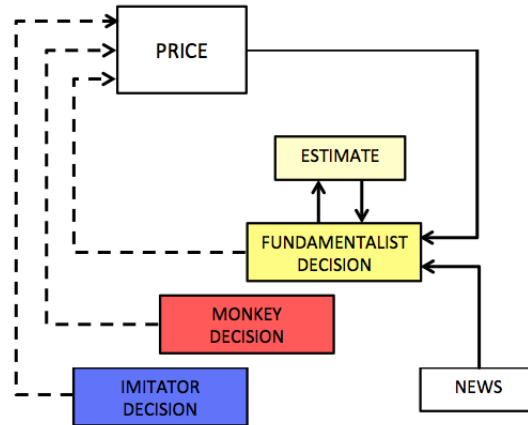
Figure 6.9: Diagram depicting the chain involving fundamentalists.
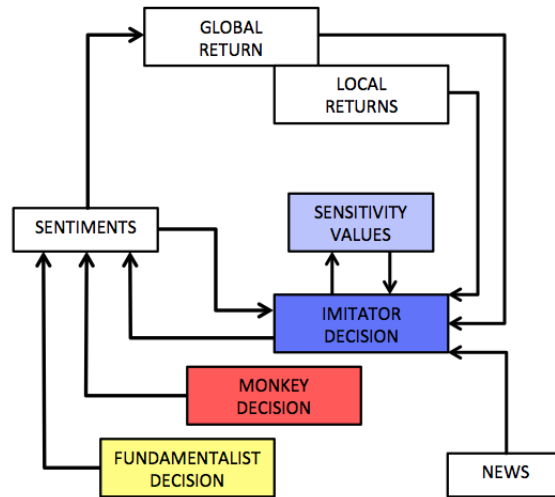
Figure 6.10: Diagram depicting the chain involving imitators.

6.11, in which the plot of the cumulative news-value was added to the plot from figure 6.6. The value of the log-price is bound by a minimum of -4.6 (roughly corresponding to a price of one cent) and a maximum of 9.2 (roughly corresponding to ten-thousand), hence the cut-off peaks.

The news is used as input by both fundamentalists and imitators. Fundamentalists adapt their estimate to the news and then compare their estimate to the current price of the share. In this procedure, the value of the news seems to have a stronger influence on the decision of the fundamentalist than the value of the price. For imitators the news makes up 1/4th of the input for their decision method. If the amount of imitators that follow the news repre-
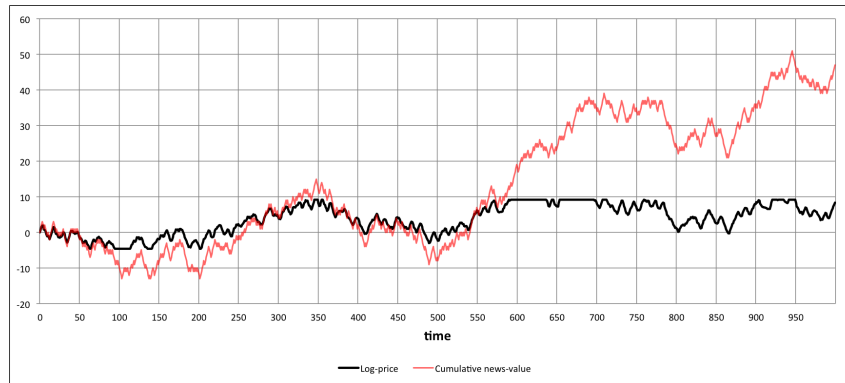
Figure 6.11: Figure 6.6 with an added plot of the cumulative news-value.

sent a majority of all traders, the market will follow the news and imitators are likely to increase their sensitivity to both the news and other imitators. This positive reinforcement explains the frequently occurring homogeneous behaviour among a majority of imitators after some time during a simulation. Imitators that deviate from the behaviour of the majority of imitators usually are traders that have not been imitators for long.

If we break the global feedback-loops of imitators by eliminating the procedure used for updating their sensitivity values, we can still observe the bubbles and crashes that follow the news. Even when there are no fundamentalists present. But when we leave the loops intact and remove the influence of the news for imitators, the bubbles and crashes are no longer observed in simulations without fundamentalists. For the fundamentalists we can do something similar. We can cut out the global feedback by letting fundamentalists compare their estimate to that of their fundamentalist neighbours instead of to the actual price. And we can eliminate the influence of the news by letting them update their estimate of the price based on the return of the previous round instead of the news. Both changes have the same effect as for the imitators. This implies that the main underlying cause of the bubbles and crashes is not a feedback-loop, but simply an external input to the system. Based on this information, we no longer consider the bubbles and crashes in the price of a share as occurring in this particular implementation of a financial market as emergent for they do not meet the requirements specified in definition 1 (see chapter 5).

When we eliminate the news as a part of the system, the price rapidly

converges to its minimum or maximum value. (See figure 6.12.) This is caused by homogeneous behaviour among both imitators and fundamentalists. I will not go into the details now, but the homogeneous behaviour of imitators is caused by the update of sensitivity-values based on global feedback. This implies that this phenomenon, although not very spectacular, *is* an example of category II emergence.
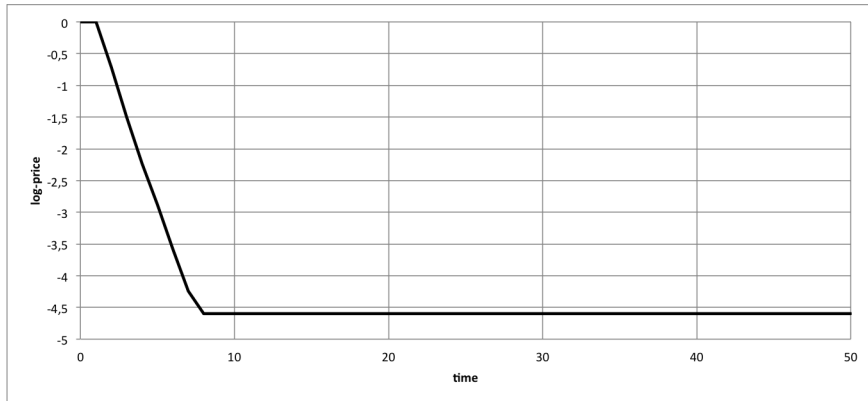
Figure 6.12: Course of the log-price when imitators are insensitive to the news and fundamentalists update their estimate of the share based on the return instead of the news, with an initial ratio of 50 % imitator and 50 % fundamentalist.

Chapter 7

Closing

This chapter concludes this thesis and is divided into two parts. I will first provide a summary of my work, answering the research questions posed at the start of this thesis and explaining its contribution to the study of emergence. I will then discuss several remaining issues concerning the concept and theory of emergence, and express my personal opinion and further ideas on the matter.

7.1   Summary

This thesis was divided into three parts: chapters 2 and 3 made up the first part, chapter 4 formed the second part, and the third part consisted of chapter 6. Chapter 5 was an intermediate chapter to summarize the findings of the first and second part before proceeding to the third part. I will now summarize the work that was done in each part of this thesis.

A lot of research went into the first part of this thesis, in which I provided an overview of emergence classifications in an attempt to clarify the concept. Despite the fact that the selection of works discussed in chapter 2 make up just a fraction of the total amount of literature on emergence, I believe those works sufficiently represent the general idea. Unlike many other works of emergence, my approach was not to add yet another definition to the already enormous accumulation of definitions, but to focus on the commonalities between several existing ones. My aim was to see whether I could unify the different emergence classifications to form a basis for a true unified theory of emergence.

I believe to have found a sufficient number of commonalities between several works comprising classifications of emergence to sucessfully construct a

unified classification of emergence based on the works I treated in chapter 2. This unified classification was worked out in section 3.3.

In the second part of this thesis I studied the underlying mechanisms of emergence. I have identified a number of elements that are required for, or play a role in, emergence in the context of multi-agent systems, by studying the different components of such systems. Several of the findings of this second part of this thesis provided additional insight and information that could be incorporated in the definitions constructed in the first part of this thesis. This was done in chapter 5, which summarized the most important findings of the first two parts of this thesis.

In the third part of this thesis I set out to determine the usability of the findings of the first and second part by studying a multi-agent system modelling an artificial financial market. I wanted to find out whether the unified classification of emergence and the set of identified underlying mechanisms of emergence could help to predict emergence in practice. In section 6.2 I have shown that the findings of the first two parts of this thesis can indeed be used to predict which types of emergence could potentially occur in the system, based on analysis of the different components of the system.

Then I analyzed two examples of actual observed phenomena in the system that could be considered emergent. Unfortunately I found that the findings of chapter 4 did not provide much to go on when analyzing the specific cause of the emergent maze-like patterns. Identifying the feedback-loops of the system pointed us towards the switch-types procedure, but this procedure was already the obvious place to look. The feedback-loops only revealed the global underlying mechanism of the emergent patterns, but did not explain why the patterns looked the way they looked. This required a mathematical analysis of the specifics in the code.

What I *did* show is that the definitions of chapter 5 can be used to verify whether the maze-like patterns were an example of category I emergence, once the underlying mechanisms that caused the patterns were identified.

The example of the bubbles and crashes occurring in the course of the price of a share was slightly different. It turned out that the bubbles and crashes occurring in this particular model were *not* emergent according to the definitions of chapter 5, because they were caused by a factor that is not part a feedback-loop.

Unfortunately, the model used for the case study in this thesis did not seem to display any category III emergence. Therefore, I have not been able to test the usability of definition 6 from chapter 5. This will be left for future work.

In conclusion, the goals set at the start of this thesis all have been fulfilled. I have successfully constructed a unified classification of emergence and inventoried important underlying mechanisms. To my knowledge, comparing and especially *joining* different definitions of emergence has not been done before. Of course I have to note that it is impossible to verify this, since the amount of literature on emergence is immense. Regardless, I believe that the main contribution to the study of emergence of the first part of this thesis is a useful and clarifying overview of the concept. As for the second part of this thesis, again, I am not aware of any work on emergence that provides an overview of the underlying mechanisms of emergence like the one in chapter 4 of this thesis.

In the following section I will elaborate on several remaining issues concerning the concept of emergence and a potential theory of emergence that I have stumbled upon during this study.

7.2 Discussion & conclusions

In chapter 6 it was shown that the findings of the first two parts of this thesis, which were captured by the definitions in chapter 5, can be used to predict whether or not emergence may occur in a system, and if so, which type of emergence will possibly be observed. The definitions of chapter 5 can furthermore be used to verify whether a phenomenon that was selected by an external observer, is emergent. That is, once the global underlying mechanisms causing it are identified. And if it is emergent, the definitions can be used to determine the category of emergence it belongs to. However, some important things are still missing.

First of all, the definitions do not provide the tools to predict or analyze the context-specific emergents that will occur in a particular system. In my opinion, the (mathematical) method employed to analyze the emergent patterns in the AFM-model in chapter 6 can not be generalized in a way that it can be used in other systems exhibiting the same form of emergence, because

it depends on the specifics of the code. The fact that the maze-like patterns turned out to result from a mistake in the code did not play a role; I have already argued that this inaccuracy did not make it any less of an example of emergence, since the patterns still resulted from feedback-loops. Besides, it turns out that in practice, many instances of emergence are 'accidents' caused by elements in the code that were overlooked during programming.

Secondly, the criteria for emergence that were specified in the definitions of chapter 5 still do not *guarantee* emergence. In chapter 5 I deliberately chose to not elaborate on the role of an external observer, which was mentioned in the general definition at the start of that chapter. The reason for this was that I believe that the main problem of emergence lies in the fact that something 'is not emergent' until it is noticed by such an external observer of the system. Once a phenomenon *is* noticed, the findings of this thesis can be used to objectively determine whether or not it is emergent. But a recurrent phenomenon, produced by feedback-loops incorporating parallel (inter-) actions of (or between) multiple elementary entities in the system, that goes unnoticed, will not be considered emergent. Recall that the verb 'to emerge' was derived from the Latin word 'emergere' which could be translated as 'to come to light'. Then one may conclude that emergence does not really listen to its own name, for it does not always emerge.

The more regularity displayed by a phenomenon meeting the emergence-criteria, the higher the chance that it will be noticed by an observer. After all, it lies in human nature to have an eye (and preference) for order, symmetry and regularity. Although, the context of a model that is being observed may play a role as well. In some rare occasions we might consider chaos to be emergent too, for instance when it deviates from the behaviour of the natural system it was supposed to model.

Some works on emergence are centered around an external observer. For instance, mathematician Edmund M. A. Ronald et al. propose a test for emergence based on whether or not an observer is surprised by a phenomenon in [32]. But others argue that this obscures the emergent phenomena. (See for instance [26].)

Subjectivity can be eliminated to some extent by using entropy as an exact method to measure the degree of disorder (or disorganization) in a system. A decrease in entropy implies an increase in self-organization and order, indi-

cating emergence. One of the works on this subject is [28], by M. Mnif and C. Müller-Schloer. This method could work for emergence such as the formation of the maze-like patterns treated in chapter 6, but it is unclear how this should work for an example such as the Game of Life. A glider somewhere in the grid probably does not have the power to decrease the overall entropy of the system. Thus, we would have to check the entropy for subsections of the grid. But, given the computational power currently available to us, it would require an unacceptable amount of time to do this for every possible subsection of a grid. Even when this would be done iteratively, for instance by starting with the whole grid and then each iteration look at partitions that are one step smaller, until something is found. However, given recent developments in quantum computing, this may become achievable one day.

In [1], computer scientist Russ Abbott proposes to use entropy as a means to verify whether an *identified* macro-level entity is emergent or not. A measure such as entropy may be used as an extra test to verify whether something is emergent in some cases, but the problem of a phenomenon having to be noticed by an external observer in the first place, remains. There are many more methods developed to formally specify and verify *known* instances of emergence, see for instance [13]. But works on *detecting* emergence are more scarce. Some examples are [29] by Niazi et al., in which proximity sensors are placed in the environment to measure flocking behaviour, and [20] by Dessalles et al., in which emergence is associated with a complexity drop. These methods seem interesting because they are exact (hence objective), but they are only suitable for detecting specific forms of emergence. This implies a need for prior knowledge concerning through which elements of the system emergence might manifest itself (e.g. through spatial distribution or through values of global variables) in order to employ such methods in the right way. And again, it is unclear how these methods could detect the emergence of a single pattern (such as the glider in the Game of Life) that is only a small part of the changes in a system.

Another issue was the deducibility-discussion that was touched upon several times throughout this thesis. Some scientists believe that there is a type of emergence of which the emergents can never be deduced to their constituents, whereas others are convinced that we will always be able to pinpoint the causes of an emergent phenomenon given that we have enough knowledge about the

system in which it occurs. The latter stems from the belief that everything has a cause. Personally, I lean more towards this idea, moreover since for the first it is unclear how we could then prove that a phenomenon is an example of that category of emergence. We would have to show that it resulted from feedback-loops while at the same time we cannot exactly pinpoint those feedback-loops, which seems contradictory. This issue seems to be discussed in [5] by Boschetti and Gray too.

When I first heard of emergence, it immediately caught my interest. What I find fascinating is that emergence produces complexity from simplicity. Something that made it even more appealing to me, is that during preparatory research for this thesis, I found that there is much discussion and disagreement on the exact meaning of the notion. This division yields great unclarity. I believe that the unification of many 'different' thoughts on emergence, that was accomplished in this thesis, takes away most of this unclarity, bringing us a step closer to understanding emergence. I furthermore believe that this unification creates a solid basis for further research on emergence.

Bibliography

[1] R. Abbott. Emergence, entities, entropy, and binding forces. In *The Agent 2004 Conference on: Social Dynamics: Interaction, Reflexivity, and Emergence*. Argonne National Labs and University of Chicago, 2004.

[2] M.A. Bedau. Weak emergence. *Philosophical Perspectives*, 11:375–399, 1997.

[3] M.A. Bedau. Downward causation and the autonomy of weak emergence. *Principia*, 6:5–50, 2002.

[4] E. Bonabeau, J.-L. Dessalles, and A. Grumbach. Characterizing emergent phenomena (1): A critical review. *Revue Internationale de Systémique*, 9(3):327–346, 1995.

[5] F. Boschetti and R. Gray. Emergence and computability. *Emergence: Complexity & Organization*, 9(1/2):120, 2007.

[6] F. Boschetti and R. Gray. A turing test for emergence. *Advances in applied self-organising systems*, 2007.

[7] M. Caldana, P. Cova, and U. Viano. Multiagent - an implementation of a multi-agent artificial financial market model in netlogo. http://ccl.northwestern.edu/netlogo/models/community/multiagent, 2006.

[8] D.T. Campbell. *Studies in the Philosophy of Biology: Reduction and Related Problems*, chapter 11: 'Downward causation' in hierarchically organized biological systems. Macmillan, London, 1974.

[9] P. Cariani. Emergence and artificial life. *Artificial Life II*, X:775–797, 1992.

[10] D. J. Chalmers. Strong and weak emergence. In *The Re-Emergence of Emergence*, pages 244–255. Oxford University Press, 2006.

[11] W. K. V. Chan, Y.-J. Son, and C. M. Macal. Agent-based simulation tutorial - simulation of emergent behavior and differences between agent-based simulation and discrete-event simulation. In B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yücesan, editors, *In proceedings of the 2010 Winter Simulation Conference*, 2010.

[12] C-C. Chen, S.B. Nagl, and C. Clack. A calculus for multi-level emergent behaviours in component-based systems and simulations. In *proceedings of Emergent Properties in Natural and Artificial Complex Systems (EP-NACS'2007), part of the 4th European Conference on Complex Systems (ECCS'07)*, pages 35–51, 2007.

[13] C-C. Chen, S.B. Nagl, and C.D. Clack. Specifying, detecting and analysing emergent behaviours in multi-level agent-based simulations. In *proceedings of the Summer Simulation Conference, Agent-directed simulation*, pages 969–976. Society for Computer Simulation International, 2007.

[14] C-C. Chen, S.B. Nagle, and C.D. Clack. Identifying multi-level emergent behaviours in agent-based simulations using complex event type specifications. *Simulation Journal special issue: Recent Advances in Unified Modeling and Simulation Approaches*, 2008, 2009.

[15] J.P. Crutchfield. The calculi of emergence: Computation, dynamics, and induction. *Physica D*, 75:11–54, 1994.

[16] P. de Grauwe and M. Grimaldi. Bubbles and crashes in a behavioural finance model. *CESifo Working Paper Series No. 1194; Riksbank Working Paper No. 164/Riksbank Research Paper Series No. 7*, 2004.

[17] J. de Haan. How emergence arises. *Ecological Complexity*, 3(4):293–301, 2006.

[18] T. W. Deacon. The hierarchic logic of emergence: untangling the interdependence of evolution and self-organization. In B. Weber and D. Depew, editors, *Evolution and learning: The baldwin effect reconsidered*. The MIT Press, 2003.

[19] J. Deguet, Y. Demazeau, and L. Magnin. Elements about the emergence issue: A survey of emergence definitions. *ComPlexUs*, 3:24–31, 2006.

[20] J.-L. Dessalles and D. Phan. Emergence in multi-agent systems: cognitive hierarchy, detection and complexity reduction. In *proceedings of the 37th European Mathematical Group Meeting*, Brest, France, June 2006.

[21] Oxford Online Dictionary. "emergence". http://oxforddictionaries.com/definition/emergence?q=emergence, 2012.

[22] M. Dorigo and L.M. Gambardella. Ant colony system : A cooperative learning approach to the traveling salesman problem. *IEEE Transactions on Evolutionary Computation*, 1(1):53–66, 1997.

[23] C. Emmeche, S. Køppe, and F. Stjernfelt. Levels, emergence, and three versions of downward causation. In N. Ole Finnemann P. Bøgh Andersen, C. Emmeche and P. Voetmann Christiansen, editors, *Downward Causation. Minds, Bodies and Matter*. Aarhus University Press, 2000.

[24] J. Fromm. Types and forms of emergence. arXiv:nlin/0506028v1, 2005.

[25] J. H. Holland. *Emergence: From Chaos to Order*. Oxford University Press, 1998.

[26] A. Kubík. Toward a formalization of emergence. *Artificial Life*, 9:41–66, 2003.

[27] Z. Li, C.H. Sim, and M.Y.H. Low. A survey of emergent behaviour and impacts in agent-based systems. In *proceedings of the 2006 IEEE International Conference on Industrial Economics*, pages 1295–1300, 2006.

[28] M. Mnif and C. Müller-Schloer. Quantitative emergence. In *proceedings of the 2006 IEEE Mountain Workshop on Adaptive and Learning Systems*, volume 49, pages 78–84. IEEE, 2006.

[29] M. Niazi and A. Hussain. Sensing emergence in complex systems. *IEEE Sensors Journal*, 11(10):2479–2480, 2011.

[30] H. Van Dyke Parunak and R. S. VanderBok. Managing emergent behavior in distributed control systems. In *proceedings of ISA-Tech '97, Instrument Society of America*, page 97, 1997.

[31] O. Paunovski, G. Eleftherakis, and T. Cowling. Framework for empirical exploration of emergence using multi-agent simulation. In *proceedings of the 2008 Workshop on Complex Systems Modelling and Simulation (CoSMoS)*, page 1. Luniver Press, 2008.

[32] E. M. A. Ronald, M. Sipper, and M. S. Capcarrère. Design, observation, surprise! a test of emergence. *Artificial Life*, 5(3):225–239, 1999.

[33] Standish, Abbass, and Bedau, editors. *Organizing Relations and Emergence*. MIT Press, 2002.

[34] H. Zhu. Formal reasoning about emergent behaviours of mas. In *proceedings of the 17th International Conference on Software Engineering and Knowledge Engineering*, pages 280–285, 2005.

Appendix A

Source code of the NetLogo model used in Chapter 6

```
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;; local variables
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

patches-own
  [ active
    ; indicates whether a trader is active (1) or not (0)
    ; (inactive means failed)

    imitator
    fundamentalist
    monkey
    ;; Indicates which type of trader it is, 0 if true, 1 if false

    type-time
    ; counts the number of successive rounds a trader has
    ; been of the current type

    random-number
    ;; It is a random number between 0 (included) and 100 (not included).

    i-sentiment
    ;; Each "imitator" can have a positive sentiment (+1),
    ;; in which case he is 'bullish', that is he believes the market will rise,
    ;; or he can have a negative sentiment (-1) in which case he is 'bearish',
    ;; that is he believes the market will fall. If the sentiment is positive the
    ;; trader buys one share if it is negative he sells one share.

    f-sentiment
    ;; It's the equivalent of i-sentiment, but for the "fundamentalists"; it is
    ;; based on rational criteria. As the name suggests, it represents the
```

80

```
;; conviction of losing or gaining money by holding the share and so
;; it determines the buy or the sell of the share.


m-sentiment
;; It's the same of the previous ones for "monkeys".


number-of-shares
;; Number of shares that each trader has (if negative it implies that the
;; trader is 'short' (we assume that there are no limits to short selling).


old-number-of-shares
;; It is the number-of-shares at time-1


liquidity
;; It represents the quantity of money which each agent has at the end
;; of every transaction.


portfolio-value
;; This variable represents each agent's portfolio, computed as
;; actual stock price times the number of stock in portfolio, after every transaction.


;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;; information used for switch-type-decision
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;


i-wealth
i-neighbors
average-i-wealth
f-wealth
f-neighbors
average-f-wealth
m-wealth
m-neighbors
average-m-wealth
;; Variables used to make a decision whether or not to change type
;; X-wealth = summed wealth of all neighbours of type X
;; X-neighbors = number of neighbors of type X
;; average-X-wealth = average wealth of all neighbors of type X



i-return
f-return
m-return
; X-return = number of optimistic neighbors (buyers) of type X
; minus the number of negative pessimistic neighbors(sellers) of type X,
; divided by the number of neighbors of type X


;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;; variables specific for imitators
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
```

```
    my-total
    ;; Represents the sum of sentiment, convinction and decision of neighbors,
    ;; weighted by propensity to be influenced by them.

    sensitivity-to-i-sentiment
    base-sensitivity-to-i-sentiment
    ; Propensity of imitators to be influenced by other imitators' sentiments
    ; regarding the news qualitative nature

    sensitivity-to-f-sentiment
    base-sensitivity-to-f-sentiment
    ; Propensity of imitators to be influenced by fundamentalists' sentiments
    ; regarding the news qualitative nature

    sensitivity-to-m-sentiment
    base-sensitivity-to-m-sentiment
    ; Propensity of imitators to be influenced by monkeys' sentiments
    ; regarding the news qualitative nature

    news-sensitivity
    ; Sensitivity that the traders have to the news qualitative meaning.

    neighbours-i-sentiments
    neighbours-f-sentiments
    neighbours-m-sentiments
    ; neighbours-X-sentiments = the summed sentiment of neighbours
    ; of type X at the end of the previous round

    ;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
    ;; variable specific for fundamentalists
    ;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

    my-estimate
    ; a fundamentalists estimate of the value of the share,
    ; based on the news

]


;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;; global variables
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

globals
  [ log-price
    log-price-\%variation
    ;; The first is the current price, due to the sum of demand
    ;; and supply; it moves according to the market returns.
    ;; The second is the variation of the price, as a percentage
    ;; of the current price, determined by demand and supply.
```

```
price
old-price
price-\%variation
;; The first is the price of the stock each time, calculated
;; as exp (log-price), the second is the same variabile at time-1,
;; while the third is the price per cent variation.

average-price
min-price
max-price
;; They measure the average, the minimum and the maximum
;; price of the stock.

average-return
min-return
max-return
;; They measure the average, the minimum and the maximum
;; return of the stock.

number-of-failed-traders
;; It counts the number of already failed agents.


average-liquidity
min-liquidity
max-liquidity
;; They measure the average, the minimum and the maximum
;; liquidity of the agents at each time.


average-portfolio-value
min-portfolio-value
max-portfolio-value
;; They measure the average, the minimum and the maximum
;; portfolio value of the agents at each time.
time
;; It measures the number of transactions in the market.

return-numerator1
return-numerator2
return-numerator3
return-numerator
return-denominator
return
;; The return of the market is calculated dividing the difference
;; between buyers and sellers by the total amount of operators.

news-qualitative-meaning
;; The news concerning the market get to every operator, and
```

```
    ;; they are the rational component of the decision

    cumulative-news-value
    ;; Cumulative value for news-qualitative-meaning

    number-of-fundamentalists
    number-of-imitators
    number-of-monkeys
    ;; They can change, because of the conversions.

    indicator-volatility
    ;; It is the variability of the market's quotations.

    init\%f
    init\%i
    init\%m
    input-traders-total
    ; User defines the ratio between different types of traders by setting the sliders.
    ; (input-traders-total is the sum over the numbers of the three sliders)
    ; This is then translated in percentages, represented by the init\%-variables above.

    global-average-f-wealth
    global-average-i-wealth
    global-average-m-wealth
    ; average wealth (portfolio-value + liquidity) per trader type

    number-changing-F-to-I
    number-changing-F-to-M
    number-changing-I-to-F
    number-changing-I-to-M
    number-changing-M-to-F
    number-changing-M-to-I
    ; number-changing-X-to-Y represents the number of traders switching
    ; from type X to type Y in the current time-step


    number-of-F+
    number-of-F-
    number-of-I+
    number-of-I-
    number-of-M+
    number-of-M-
  ]


 ;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
 ;; setup
 ;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

to setup
  ca
```

```
set log-price 0
set price exp log-price
set time 0
set cumulative-news-value 0

; set general initial values for traders
ask patches
[ reset-type-labels
  set active 1
  set liquidity endowment
  set number-of-shares 1
]


; translate user input to percentages regarding the distribution
; of agents over 3 types
set input-traders-total (fundamentalists + imitators + monkeys)
ifelse input-traders-total = 0
[ user-message "Number of traders = zero, please set at least one of
  the trader-type-sliders to non-zero." stop ]
[ set init%f (100 / input-traders-total) * fundamentalists
  set init%i (100 / input-traders-total) * imitators
  set init%m (100 / input-traders-total) * monkeys ]


; divide traders over different types
; random-number is a number between 0 and 100
; if random-number is below init%f, patch becomes a fundamentalist
; if random-number is between init%f and (init%f + init%i),
; patch becomes an imitator
; if random-number is between (init%f + init%i) and 100,
; patch becomes a monkey
ask patches
[ set random-number random-float 100
  ifelse random-number < init%f
    [ set fundamentalist 1
      set my-estimate 0
      set pcolor yellow ]  ; fundamentalists are yellow
    [ ifelse random-number < (init%f + init%i)
      [ set imitator 1
        set pcolor blue ]  ; imitators are blue
      [ set monkey 1
        set pcolor red ]   ; monkeys are red
    ]
]


; set initial values for imitators
ask patches with [ imitator = 1 ]
[ set news-sensitivity (random-float max-news-sensitivity)
  set base-sensitivity-to-i-sentiment
    (random-float max-base-sensitivity-to-i-sentiment)
  set sensitivity-to-i-sentiment
```

```
      base-sensitivity-to-i-sentiment  ;;propensity to copy other i's
    set base-sensitivity-to-f-sentiment
      (random-float max-base-sensitivity-to-f-sentiment)
    set sensitivity-to-f-sentiment
      random-float max-base-sensitivity-to-f-sentiment ;;propensity to copy f's
    set base-sensitivity-to-m-sentiment
      (random-float max-base-sensitivity-to-m-sentiment)
    set sensitivity-to-m-sentiment
      random-float max-base-sensitivity-to-m-sentiment ;;propensity to copy m's
  ]


  ; count the number of traders for each type
  set number-of-fundamentalists count patches with [ fundamentalist = 1 ]
  set number-of-imitators count patches with [ imitator = 1 ]
  set number-of-monkeys count patches with [ monkey = 1 ]


  ; plot the initial values
  do-plot
end


;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;; each round
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;


to go
 set time time + 1
 ; new round, so increase time
 reset-change-counters
 ; reset all counters for monitoring the number of type-changes at the start of the round
 old-number
 ; update the number of shares (current number becomes old number since we started a new round)
 switch-types
 ; determine for each trader whether or not it will switch types
 news-arrival
 ; determine the news-value for the current round
 imitators-decision
 ; determine for each imitator whether he will buy or sell an asset
 fundamentalists-decision
 ; determine for each fundamentalist whether he will buy or sell an asset
 monkeys-decision
 ; determine for each monkey whether he will buy or sell an asset
 execute-decisions
 ; update colours, sentiments, liquidity and portfolio-values
 update-return-and-price
 ; calculate the return and update the price
 fail-or-survive
 ; check whether there are traders that have become in too much debt
 update-neighbour-info
 ; update the info about your neighbours
 update-imitators-sensitivity-values
```

```
 ; update the sensitivity-values of imitators
 update-info-for-monitoring
 ; update info used for analysis of the system (plots in the interface)
 do-plot
 ; plot the graphs in the interface
end




;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;; Decision methods for trading
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;


 ;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
 ;; decision-method for imitators
 ;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
 ;; Whether the "imitators" are optimist and buy the share,
 ;; depends on the four elements you can find here.
 ;; If the sum of them is positive the "imitators" are optimist.

 to imitators-decision
 ask patches with [ imitator = 1 ]
   [ set my-total ( sensitivity-to-i-sentiment * i-return +
                    sensitivity-to-f-sentiment * f-return +
                    sensitivity-to-m-sentiment * m-return +
                    news-sensitivity * (news-qualitative-meaning) )
   ]



 ask patches with [ imitator = 1 ]
   [ ifelse (my-total > 0)
     [ set number-of-shares (number-of-shares + 1) ]
     [ ifelse (my-total < 0) ; my-total is smaller than zero, yielding a pessimistic mood
       [ set number-of-shares (number-of-shares - 1) ]
       [ ifelse (random-float 1.0 >= 0.5) ; my-total = zero, decide randomly
           [ set number-of-shares (number-of-shares + 1) ]
           [ set number-of-shares (number-of-shares - 1) ]
       ]
     ]
   ]


 end



;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; fundamentalist's decision
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;; The "fundamentalist" buy the share if they think it is underestimated,
```

```
;; otherwise they sell.

to fundamentalists-decision
 ask patches with [ fundamentalist = 1 ]
 [ ifelse (news-qualitative-meaning = 1)
    [ set my-estimate my-estimate + random-float 1.00001
      if (my-estimate > 10.2)
        [ set my-estimate 10.2 ]
   ]
    [ set my-estimate my-estimate - random-float 1.00001
      if (my-estimate < -5.6)
        [ set my-estimate -5.6 ]
   ]
 ]

 ask patches with [ fundamentalist = 1 ]
 [ ifelse (my-estimate > log-price)
    [ set number-of-shares number-of-shares + 1 ]
    [ set number-of-shares number-of-shares - 1 ]
 ]
end

;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; monkey's decision
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;; The "monkeys" buy or sell the share randomly.

to monkeys-decision
  ask patches with [ monkey = 1 ]
  [ ifelse (random-float 1.00) >= .5
    [ set number-of-shares number-of-shares + 1 ]
    [ set number-of-shares number-of-shares - 1 ]
  ]
  end



;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;; Decision-method for switching types
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; Every trader calculates the average wealth of each trader type
; among his active neighbours. If the average wealth of the type
; with the highest average wealth is over 'change-threshold' percent
; higher than his own current wealth, the trader will switch to that
; type with a chance of 'change-probability' percent
; (given that it is not the same type as his current type)

to switch-types
  ask patches with [ active = 1 ]
  [  if ( ((portfolio-value + liquidity) +
     ((portfolio-value + liquidity) * change-threshold)) < average-f-wealth or
```

```
      ((portfolio-value + liquidity) +
((portfolio-value + liquidity) * change-threshold)) < average-i-wealth or
  ((portfolio-value + liquidity) +
      ((portfolio-value + liquidity) * change-threshold)) < average-m-wealth)
  [ ; determine wealthiest trader type among neighbors
    ifelse (average-f-wealth >= average-m-wealth) and
    (average-f-wealth >= average-i-wealth)
      [ ; fundamentalists are wealthiest
        ; If you are not already a fundamentalist,
        ; then become one with chance 'change-probability'
        if fundamentalist != 1 and (random-float 1.00 < change-probability)
          [ if imitator = 1 [ set number-changing-I-to-F (number-changing-I-to-F + 1)
                              clear-i-vars ]
            if monkey = 1 [ set number-changing-M-to-F (number-changing-M-to-F + 1) ]
            reset-type-labels
            set fundamentalist 1
            set my-estimate log-price
            set type-time 0
          ]
      ]
      [ ifelse (average-i-wealth >= average-f-wealth) and (average-i-wealth >= average-m-wealth)
          [ ; imitators are wealthiest
            ; If you are not already a imitator, then become one with chance 'change-probability'
            if imitator != 1 and (random-float 1.00 < change-probability)
              [ if fundamentalist = 1 [ set number-changing-F-to-I (number-changing-F-to-I + 1)
                                        set my-estimate 0 ]
                if monkey = 1 [ set number-changing-M-to-I (number-changing-M-to-I + 1) ]
                reset-type-labels
                set imitator 1
                set news-sensitivity (random-float max-news-sensitivity)
                set base-sensitivity-to-i-sentiment
                  (random-float max-base-sensitivity-to-i-sentiment)
                set sensitivity-to-i-sentiment base-sensitivity-to-i-sentiment
                set base-sensitivity-to-f-sentiment
                  (random-float max-base-sensitivity-to-f-sentiment)
                set sensitivity-to-f-sentiment
                  random-float max-base-sensitivity-to-f-sentiment
                set base-sensitivity-to-m-sentiment
                  (random-float max-base-sensitivity-to-m-sentiment)
                set sensitivity-to-m-sentiment
                  random-float max-base-sensitivity-to-m-sentiment
                set type-time 0
              ]
          ]
          [ ;monkeys are wealthiest
            ;If you are not already a monkey, then become one with chance 'change-probability'
            if monkey != 1 and (random-float 1.00 < change-probability)
              [ if fundamentalist = 1 [ set number-changing-F-to-M (number-changing-F-to-M + 1) ]
                if imitator = 1 [ set number-changing-I-to-M (number-changing-I-to-M + 1)
                                  clear-i-vars ]
```

```
                        reset-type-labels
                        set monkey 1
                        set type-time 0
                    ]
                ]
            ]
        ]
        set type-time (type-time + 1) ;increase the counter keeping the time a trader stays the same type
    ]
 end
```

```
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;; Update-methods
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;


;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; News Arrival mechanism
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;; The news act as a random normal variable and they change the value of the share.

to news-arrival
 ifelse (random-float 1) >= .5
   [ set news-qualitative-meaning 1 ]
   [ set news-qualitative-meaning -1 ]
 set cumulative-news-value cumulative-news-value + news-qualitative-meaning
end


;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; Number of shares at time-1
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;; This updates the variable "old-number-of-shares" to the number
;; of shares which every agent had in her portfolio at time-1

to old-number
ask patches
[set old-number-of-shares number-of-shares]
end


;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; Execute the decisions of traders
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; Sentiments and colours are updated according to the decisions of traders
; Shares are sold and bought for the current price

to execute-decisions
  ask patches with [ active = 1 ]
  [ reset-sentiments
```

90

```
      set portfolio-value (price * number-of-shares)
      ifelse number-of-shares > old-number-of-shares
        [ set liquidity liquidity - price  ; bought a share, so pay for it
          ifelse fundamentalist = 1
            [ set f-sentiment 1
              set pcolor yellow ]
            [ ifelse imitator = 1
              [ set i-sentiment 1
                set pcolor blue ]
              [ set m-sentiment 1
                set pcolor red ]
          ]
      ]
      [ set liquidity liquidity + price  ; sold a share
          ifelse fundamentalist = 1
            [ set f-sentiment -1
              set pcolor yellow + 3 ]
            [ ifelse imitator = 1
              [ set i-sentiment -1
                set pcolor blue + 1 ]
              [ set m-sentiment -1
                set pcolor red + 1 ]
          ]
      ]
  ]
end


;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; Calulate return and price
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; We compute the price and the return, as they result from the actions
;  undertaken by the operators. The return is the excess of demand over
; the supply, divided by the dimension of the market.


to update-return-and-price
  ;; The denominator is the number of active traders
  set return-denominator count patches with [ active = 1 ]
  ;; The numerator is the difference between sellers and buyers
  set return-numerator1 sum [f-sentiment] of patches
  set return-numerator2 sum [i-sentiment] of patches
  set return-numerator3 sum [m-sentiment] of patches
  set return-numerator
    (return-numerator1 + return-numerator2 + return-numerator3)
  ;; The return modifies the price of the share (return = average sentiment of the market)

  ask patches
  [ ifelse return-denominator = 0
    [ set return 0 show "no active traders left" stop]
    [ set return return-numerator / return-denominator ]
  ]
```

```
    set old-price price
    set log-price log-price + return
    if log-price < -4.6 [ set log-price -4.6 ]
    ; set minimum log-price so that minimum price per share =  0.01 (1 cent)
    if log-price > 9.2 [ set log-price 9.2 ]
    ; set maximum log-price so that maximum price per share =  10,000
    set price exp log-price
end



;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; Check whether traders have failed
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; If a traders' wealth is below the maximum debt, he is excluded from the market
; If only his liquidity is below the maximum debt, he has to sell shares until
; his debt is at an acceptable level
; If he does not have enough shares to accomplish this, he will be excluded from the market
; If only his portfolio-value is below the maximum debt, he has to buy shares
; until his debt is at an acceptable level
; If he does not have enough money (liquidity) to accomplish this, he will
; be excluded from the market

to fail-or-survive
  ask patches with [ active = 1 ]
  [ ifelse (liquidity + portfolio-value) < maximum-debt
    [ set pcolor black
      reset-type-labels
      set active 0
      set number-of-shares 0
      set f-sentiment 0
      set m-sentiment 0
      set i-sentiment 0
    ]
    [ ifelse liquidity < maximum-debt and (liquidity + portfolio-value) >= maximum-debt
      [ while [ liquidity < maximum-debt and (liquidity + portfolio-value) >= maximum-debt ]
        [ set number-of-shares (number-of-shares - 1)
          set portfolio-value (price * number-of-shares)
          set liquidity liquidity + price
        ]
      ]
    [ if portfolio-value < maximum-debt and (liquidity + portfolio-value) >= maximum-debt
      [ while [ portfolio-value < maximum-debt and
        (liquidity + portfolio-value) >= maximum-debt ]
        [ set number-of-shares (number-of-shares + 1)
          set liquidity liquidity - price
          set portfolio-value (price * number-of-shares) ]
      ]
    ]
  ]
```

```
]

  ; for traders that had to sell or buy shares in order to pay their debt:
  ; if you did not have enough resources to get out of debt, you will
  ; be excluded from the market
  ask patches with [ active = 1 ]
    [ if (liquidity + portfolio-value) < maximum-debt
      [ set pcolor black
        reset-type-labels
        set active 0
        set number-of-shares 0
        set f-sentiment 0
        set m-sentiment 0
        set i-sentiment 0
      ]
    ]


set number-of-failed-traders count patches with [ active = 0 ]

ifelse ((count patches with [ active = 1 ]) != 0)
  [ set average-portfolio-value (sum [ portfolio-value ] of
    patches with [ active = 1 ]) / (count patches with [ active = 1 ])
    set average-liquidity
      (sum [ liquidity ] of patches with [ active = 1 ] ) / (count patches with [ active = 1 ]) ]
  [ set average-portfolio-value 0
    set average-liquidity 0 ]

set min-portfolio-value min [ portfolio-value ] of patches with [ active = 1 ]
set max-portfolio-value max [ portfolio-value ] of patches with [ active = 1 ]
set min-liquidity min [liquidity] of patches
set max-liquidity max [liquidity] of patches
end


;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; Update information about neighbours
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; Count the number of neighbours of each type, their sentiments,
; their return and average wealth
to update-neighbour-info
  ask patches with [ active = 1 ]
  [ ; count the number of neighbours per trader type
    set i-neighbors sum [ imitator ] of neighbors
    set f-neighbors sum [ fundamentalist ] of neighbors
    set m-neighbors sum [ monkey ] of neighbors

    ; sum the sentiments of neighbours per trader type
    set neighbours-i-sentiments sum [ i-sentiment ] of neighbors
    set neighbours-f-sentiments sum [ f-sentiment ] of neighbors
    set neighbours-m-sentiments sum [ m-sentiment ] of neighbors
```

```
    ; calculate the return of neighbours per trader type
    ifelse (i-neighbors != 0)
        [ set i-return ( neighbours-i-sentiments / i-neighbors ) ]
        [ set i-return 0]
    ifelse (f-neighbors != 0)
        [ set f-return ( neighbours-f-sentiments / f-neighbors ) ]
        [ set f-return 0 ]
    ifelse (m-neighbors != 0)
        [ set m-return ( neighbours-m-sentiments / m-neighbors ) ]
        [ set m-return 0 ]

    ; calculate the 'wealth' of neighbours per trader type
    set f-wealth
      f-wealth + sum [portfolio-value + liquidity] of neighbors with [ fundamentalist = 1 ]
    set i-wealth
      i-wealth + sum [portfolio-value + liquidity] of neighbors with [ imitator = 1 ]
    set m-wealth
      m-wealth + sum [portfolio-value + liquidity] of neighbors with [ monkey = 1 ]

    ; calculate the average wealth of neighbours per trader type
    ifelse f-neighbors != 0
      [ set average-f-wealth (f-wealth / f-neighbors) ]
      [ set average-f-wealth 0 ]
    ifelse i-neighbors != 0
      [ set average-i-wealth (i-wealth / i-neighbors) ]
      [ set average-i-wealth 0 ]
    ifelse m-neighbors != 0
      [ set average-m-wealth (m-wealth / m-neighbors) ]
      [ set average-m-wealth 0 ]
  ]
end


;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; Update sensitivity values of imitators
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
; Calculate the return of each trader type among your neighbours
; Update the sensitivity-values as follows:
; if the price increased (return > 0) and the average sentiment
;     of neighbors of type X was positive,
;         then sensitivity-to-X-sentiment is increased, otherwise it is decreased
; if the price decreased (return < 0) and the average sentiment of
;     neighbors of type X was negative,
;         then sensitivity-to-X-sentiment is increased, otherwise it is decreased
; if the price stayed the same (return = 0), sensitivity-to-sentiments stays the same too
; the news-sensitivity is increased by the absolute return when the movement of the market
; corresponds to the value of the news, and decreased by the return otherwise

to update-imitators-sensitivity-values
  ask patches with [ imitator = 1 ]
    [ if (return > 0)
```

```
        [ set sensitivity-to-i-sentiment  base-sensitivity-to-i-sentiment + i-return
          set sensitivity-to-f-sentiment  base-sensitivity-to-f-sentiment + f-return
          set sensitivity-to-m-sentiment  base-sensitivity-to-m-sentiment + m-return
          ifelse (news-qualitative-meaning > 0)
          ; if the movement of the market corresponded to the value of the news
            [ set news-sensitivity news-sensitivity + return ; increase the sensitivity to the news
              if (news-sensitivity > max-news-sensitivity)
                [ set news-sensitivity max-news-sensitivity ]
            ]
            [ set news-sensitivity news-sensitivity - return ; otherwise, decrease the sensitivity
              if (news-sensitivity < 0)
                [ set news-sensitivity 0 ]
            ]
          ]
      if (return < 0)
        [ set sensitivity-to-i-sentiment  base-sensitivity-to-i-sentiment - i-return
          set sensitivity-to-f-sentiment  base-sensitivity-to-f-sentiment - f-return
          set sensitivity-to-m-sentiment  base-sensitivity-to-m-sentiment - m-return
          ifelse (news-qualitative-meaning < 0)
          ; if the movement of the market corresponded to the value of the news
            [ set news-sensitivity news-sensitivity - return ; increase the sensitivity to the news
              if (news-sensitivity > max-news-sensitivity)
                [ set news-sensitivity max-news-sensitivity ]
            ]
            [ set news-sensitivity news-sensitivity + return ; otherwise, decrease the sensitivity
              if (news-sensitivity < 0)
                [ set news-sensitivity 0 ]
            ]
          ]

    ;keep sensitivity-values between 0 and 1
    if sensitivity-to-i-sentiment < 0 [ set sensitivity-to-i-sentiment 0 ]
    if sensitivity-to-i-sentiment > 1 [ set sensitivity-to-i-sentiment 1 ]
    if sensitivity-to-f-sentiment < 0 [ set sensitivity-to-f-sentiment 0 ]
    if sensitivity-to-f-sentiment > 1 [ set sensitivity-to-f-sentiment 1 ]
    if sensitivity-to-m-sentiment < 0 [ set sensitivity-to-m-sentiment 0 ]
    if sensitivity-to-m-sentiment > 1 [ set sensitivity-to-m-sentiment 1 ]
  ]
end



;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;; Procedures for resetting variable-values
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

; to ensure a trader cannot be of multiple types at the same time
to reset-type-labels
  set fundamentalist 0
  set imitator 0
  set monkey 0
```

```
 end

 ; clean up the imitator-specific variables when a trader switches to another type
 to clear-i-vars
   set news-sensitivity 0
   set base-sensitivity-to-i-sentiment 0
   set sensitivity-to-i-sentiment 0
   set base-sensitivity-to-f-sentiment 0
   set sensitivity-to-f-sentiment 0
   set base-sensitivity-to-m-sentiment 0
   set sensitivity-to-m-sentiment 0
   set my-total 0
 end


; ensure that traders that have changed types do not keep their old sentiment-value
to reset-sentiments
  set f-sentiment 0
  set i-sentiment 0
  set m-sentiment 0
end

; counting the number of traders changing between types
; has to be reset at the start of each round
to reset-change-counters
  set number-changing-F-to-I 0
  set number-changing-F-to-M 0
  set number-changing-I-to-F 0
  set number-changing-I-to-M 0
  set number-changing-M-to-F 0
  set number-changing-M-to-I 0
end

 ;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
 ;; Procedures for keeping information during simulation for observational purposes
 ;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;

 to update-info-for-monitoring

  ; calculate volatility in price
  ; (volatility of the share market refers to the absolute value of the return)
  ; set indicator-volatility abs(return)

  ; update variation in price
  ifelse time = 1
     [ set price-\%variation 0 ]
     [ set price-\%variation ln (price / old-price) ]

   ; set avg, min & max-values for price & return for time 1
   if time = 1
```

```
  [ set average-price price
    set min-price price
    set max-price price
    set average-return return
    set min-return return
    set max-return return ]


; calculate the current average price & return
set average-price average-price * (time - 1) / time + price / time
set average-return average-return * (time - 1) / time + return / time


; update the min & max-values for price & return if necessary
if (price < min-price) [ set min-price price ]
if (price > max-price) [ set max-price price ]
if (return < min-return) [ set min-return return ]
if (return > max-return) [ set max-return return ]


; count the number of traders for each type
set number-of-fundamentalists count patches with [ fundamentalist = 1 ]
set number-of-imitators count patches with [ imitator = 1 ]
set number-of-monkeys count patches with [ monkey = 1 ]


; count the number of buyers (optimists) and sellers (pessimists) for each type
set number-of-F+ count patches with [ fundamentalist = 1 and f-sentiment = 1 ]
set number-of-F- count patches with [ fundamentalist = 1 and f-sentiment = -1 ]
set number-of-I+ count patches with [ imitator = 1 and i-sentiment = 1 ]
set number-of-I- count patches with [ imitator = 1 and i-sentiment = -1 ]
set number-of-M+ count patches with [ monkey = 1 and m-sentiment = 1 ]
set number-of-M- count patches with [ monkey = 1 and m-sentiment = -1 ]


; calculate the average wealth of all traders of each type
ifelse number-of-fundamentalists != 0
  [ set global-average-f-wealth ( (sum [portfolio-value + liquidity] of
    patches with [ fundamentalist = 1 ]) / number-of-fundamentalists ) ]
  [ set global-average-f-wealth 0 ]
ifelse number-of-imitators != 0
  [ set global-average-i-wealth ( (sum [portfolio-value + liquidity] of
    patches with [ imitator = 1 ]) / number-of-imitators ) ]
  [ set global-average-i-wealth 0 ]
ifelse number-of-monkeys != 0
  [ set global-average-m-wealth ( (sum [portfolio-value + liquidity] of
    patches with [ monkey = 1 ]) / number-of-monkeys ) ]
  [ set global-average-m-wealth 0 ]

end


;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
;; Plot the graphs
;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;;
```

```
to do-plot
 set-current-plot "Price"
 set-current-plot-pen "price"
 plot log-price

 set-current-plot "Return"
 set-current-plot-pen "return"
 plot return

 set-current-plot "Cumulative news-value"
 set-current-plot-pen "value"
 plot cumulative-news-value
 set-current-plot-pen "zero"
 plot 0

 set-current-plot "Volatility"
 set-current-plot-pen "indicator-volatility"
 plot indicator-volatility

 set-current-plot "Price-%variation"
 set-current-plot-pen "Price-%variation"
 plot price-%variation

 set-current-plot "Portfolio Value"
 set-current-plot-pen "Min"
 plot min-portfolio-value
 set-current-plot-pen "Max"
 plot max-portfolio-value
 set-current-plot-pen "Avg"
 plot average-portfolio-value

 set-current-plot "Liquidity"
 set-current-plot-pen "Min"
 plot min-liquidity
 set-current-plot-pen "Max"
 plot max-liquidity
 set-current-plot-pen "Avg"
 plot average-liquidity

 set-current-plot "Distribution of types"
 set-current-plot-pen "x"
 plot number-of-failed-traders
 set-current-plot-pen "f"
 plot number-of-fundamentalists
 set-current-plot-pen "I"
 plot number-of-imitators
 set-current-plot-pen "M"
 plot number-of-monkeys

 set-current-plot "Average wealth per type"
```

```
  set-current-plot-pen "f"
  plot global-average-f-wealth
  set-current-plot-pen "i"
  plot global-average-i-wealth
  set-current-plot-pen "m"
  plot global-average-m-wealth

  set-current-plot "Type-changes for fundamentalists"
  set-current-plot-pen "F-to-I"
  plot number-changing-F-to-I
  set-current-plot-pen "F-to-M"
  plot number-changing-F-to-M

  set-current-plot "Type-changes for Imitators"
  set-current-plot-pen "I-to-F"
  plot number-changing-I-to-F
  set-current-plot-pen "I-to-M"
  plot number-changing-I-to-M

  set-current-plot "Type-changes for monkeys"
  set-current-plot-pen "M-to-F"
  plot number-changing-M-to-F
  set-current-plot-pen "M-to-I"
  plot number-changing-M-to-I

  set-current-plot "Distribution of f-sentiments"
  set-current-plot-pen "F+"
  plot number-of-F+
  set-current-plot-pen "F-"
  plot number-of-F-

  set-current-plot "Distribution of i-sentiments"
  set-current-plot-pen "I+"
  plot number-of-I+
  set-current-plot-pen "I-"
  plot number-of-I-

  set-current-plot "Distribution of m-sentiments"
  set-current-plot-pen "M+"
  plot number-of-M+
  set-current-plot-pen "M-"
  plot number-of-M-
end
```