FAKULTÄT FÜR !NFORMATIK
Faculty of Informatics

Innovation - The Emergence of Novelty

DIPLOMARBEIT

zur Erlangung des akademischen Grades

Diplom-Ingenieur

im Rahmen des Studiums

Wirtschaftsinformatik

eingereicht von

Ing. Markus Rohringer, BSc

Matrikelnummer 0627776

an der
Fakultät für Informatik der Technischen Universität Wien

Betreuung: Univ.Prof. Mag. Dr. Hardy Hanappi

Wien, 18.02.2015

(Unterschrift Verfasser) (Unterschrift Betreuung)

FAKULTÄT FÜR !NFORMATIK
Faculty of Informatics

Innovation - The Emergence of Novelty

MASTER'S THESIS

submitted in partial fulfillment of the requirements for the degree of

Diplom-Ingenieur

in

Business Informatics

by

Ing. Markus Rohringer, BSc

Registration Number 0627776

to the Faculty of Informatics
at the Vienna University of Technology

Advisor: Univ.Prof. Mag. Dr. Hardy Hanappi

Vienna, 18.02.2015

(Signature of Author) (Signature of Advisor)

Erklärung zur Verfassung der Arbeit

Ing. Markus Rohringer, BSc
Hausfeldstrasse 21/10, 1220 Wien

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit - einschließlich Tabellen, Karten und Abbildungen -, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.

(Ort, Datum) (Unterschrift Verfasser)

Abstract

The rather unexpected emergence of the still ongoing economic crisis has shown that not only the economic system itself has its flaws, but also that the equilibria-based models to describe it are insufficient. Evolutionary economics is an alternative approach that allows (economic) evolution and revolution in the development of its models and is therefore believed to describe the real world much more appropriate. However, there are still a lot of open questions on how to model the emergence of novelty - new elements that lead to progress - in such frameworks. Innovation is an undeniable force regarding economic growth, but it is obviously impossible to predict which concrete innovations will be implemented in the future. What can be done though is to create a general mechanism within the model that is able to generate new abstract elements which then again influence the model itself. One possible way of doing this, as it is often the case in real world innovations, is using old elements and recombine them to create new ones. The agent-based model that will be implemented with this work will use this approach to demonstrate how it can be utilized to define how producers choose their new products, how they recombine production processes and how they can test their products on the market.

Contents

CHAPTER 1

Introduction

1.1 Motivation & Problem Statement

The rather unexpected emergence of the still ongoing economic crisis has shown that not only the economic system itself has its flaws, but also that the equilibria-based models to describe it are insufficient. *Evolutionary economics* is an alternative approach that allows (economic) evolution and revolution in the development of its models and is therefore believed to describe the real world much more appropriate. However, there are still a lot of open questions on how to model the *emergence of novelty* - new elements that lead to progress - in such frameworks.

Innovation is an undeniable force regarding economic growth, but it is obviously impossible to predict which concrete innovations will be implemented in the future. What can be done though is to create a general mechanism within the model that is able to generate new abstract elements which then again influence the model itself. One possible way of doing this, as it is often the case in real world innovations, is using old elements and *recombine* them to create new ones. The *agent-based model* that will be implemented with this work will use this approach to demonstrate how it can be utilized to define how producers choose their new products, how they recombine production processes and how they can test their products on the market.

1.2 Aim of the Work

The aim of this work is to explore and show how the emergence of novelty can be implemented in an evolutionary economic framework - by exactly doing this in an concrete agent-based simulation. While the implementation of the model and the exploration of the simulation results is the main focus of this work for sure, the theoretical part is important too. Evolutionary economics is still a quite new and growing field without strict boundaries, which is why a lot of the work done there is relatively scattered and unrelated. Starting with Schumpeter and his idea of innovation and his famous concept of the *entrepreneur*, this work will try to tie these different strings related to the emergence of novelty. Modern ideas and concepts like those of recent authorities in the field, e.g. Witt, Foster, Hodgson or Dopfer, will play an important role.

As already mentioned the aim of the model itself is mainly to build a showcase on how a concrete model can use those theoretical ideas. The model built in this work will focus on the production side, while using the market only as a sort of testing mechanism, to find out how good the new products fit the customers' demands. The innovator (or entrepreneur) agents have access to a pool (which has not to be the same for every agent) of (partial) production processes. For those processes certain parameters are known from historical experience, like input (costs), quantitative output and the output in terms of customer utility. This utility is characterized by a vector of quantitative measures for different product properties which can be matched against a similar vector which represents the customer's expectations. The innovators will try to recombine those production processes to create a new process which represents a new product.

A key concept for this model are therefore the links between those processes, which define if a combination is possible at all, and if so, what the already mentioned parameters of the new process look like. When leaving the view of a single agent and zooming out to an aggregated view, this will lead to a dynamic race for innovation, where older products will experience an utility discount over time and the different producers will try to come up with new products, do constant market research and eventually produce successfully tested new products, while being in permanent competition with each other. The questions to be answered are: What information does the simulation provide regarding the emergence of novelty in such an environment? How well do the simulation results match with what we can experience in reality? How can this model be enhanced or used in other more sophisticated evolutionary models, that need such an innovation engine, but focus on other key topics?

1.3 Methodological Approach

The methodological approach consists of the following steps:

1. **Literature Review**

   Existing research and literature will be gathered, studied, merged and used as theoretical basis for the model.

2. **Agent-based & Network Modelling**

   The agent-based model to represent the innovation system must be designed. It will consist of a market as an utility testing system and producers with access to pools of production processes, which are themselves re-combinable and therefore linked to each other in some kind of network, which has to be modeled as well.

3. **Agent-based Simulation**

The conceptual model developed in the previous step has to be implemented in an agent-based modeling and simulation environment. This will be done in NetLogo. After a successful model verification various simulation runs will be performed to explore the behavior space and the validity of the model. The results of these experiments and the usability of the model for future applications will be analyzed and interpreted from an evolutionary economic point of view.

1.4 Structure of the Work

This work follows a structure that the author thinks is the most useful to approach this topic. After this introduction the first step will be to look at the general treatment of the topic of innovation and growth in economics, which will be done in the second chapter. A logical starting point for doing this is to give a brief overview over traditional growth models. Afterwards we will continue with one of the most influential concepts regarding innovation in economics, which is Schumpeter's entrepreneur. To conclude this chapter we will examine what modern mainstream has to offer in terms of innovation models.

In chapter 3 we will approach the theoretical backbone of this work, which is evolutionary economics. We will see, that even among its advocates, there is not an unified perspective about what evolutionary economics really is or what it should be. But we will try to find the common ground as well as to highlight the differences. Afterwards we will get more specific and will examine the importance of novelty in this evolutionary context, because as we will see, evolution without novelty is not possible. But of course we will also have to look what different things can be novel. This automatically leads us to the last part of this chapter, where we will outline how novelty can be modeled, but of course with an emphasis on agent-based modeling, because this is more or less already the introduction to the model presented in this work.

Chapter 4 will be the centerpiece of this work insofar as it describes the uniqueness of it. The model which is introduced with this work will be presented and explained here. Furthermore the *NetLogo* implementation will be described and the results of the simulation runs will be first described and explored and finally also economically interpreted.

The work will end in chapter 5 with a general conclusion about the theoretical foundations of this model, implications of the results and of course also how the model can be used in further research and which possible improvements are imaginable.

Below you can see the exact structure as it was described now:

1. **Introduction**
2. **Innovation and Growth in Economics**
    - Traditional Growth Models

- Schumpeter and his Entrepreneur
- Modern Mainstream Innovation Models

3. **Evolutionary Economics and the Emergence of Novelty**
   - Evolutionary Economics in a Nutshell
   - The Importance of the Emergence of Novelty in an Evolutionary Context
   - How to model Novelty
4. **Modeling and Simulation**
   - A Model of Innovation with re-combinable Production Processes
   - Agent-based Implementation with NetLogo
   - Simulation Runs and Exploration of the Behavior Space
   - Economic Interpretation of the Results
5. **Conclusion**

1.5 State of the Art

When considering the state of the art, one has to distinguish between *neoclassical* and evolutionary economic theory, as the latter unfortunately hardly can be called state of the art in mainstream economics. Traditional analytical growth models almost ignored innovation as they were not able to incorporate it. This was most impressively demonstrated by Solow (1957), who was forced to handle technological progress as an exogenous process. The so-called *Solow residual* is the part of economic growth that can not be attributed to accumulation of inputs. For the United States, it represented 80 percent of the growth that could not be explained by the typical rational and cost minimizing behavior of the neoclassical world.

Of course neoclassical growth models have been advanced since then. Those models explicitly incorporate intentional innovation in terms of expanding product variety and rising product quality. They do this by introducing knowledge capital, brand proliferation, industrial policies and so on. However, these are still analytical equilibrium models. In the best case they describe how innovation influences the economy but they can not explain the emergence of novelty at all. Neither the behavioral reasons of innovators to innovate nor the way, in which new things can arise.

In evolutionary economics on the other hand, these questions are elementary and therefore oft-treated topics. The most common method for the emergence of novelty, which is also used in this work, is recombination. Old elements are recombined to create new elements. The idea itself is obviously not very new, in fact already Schumpeter characterized it. However, to my knowledge, most of the work done in this regard is of theoretical, often even very epistemological nature. There are suggestions on how to model this method (Witt [25], Hanappi and Hanappi-Egger

[14]), but still in a very general way. Using this recombination idea for production processes with a network of links that define the inputs and outputs of those new combinations should be quite novel. The same is true for representing the utility of resulting products via vectors of quantitative measures and matching them against corresponding customer expectation vectors. Also, applying this idea (and also other evolutionary economic concepts) to a concrete agent-based simulation is something that should be done more often in the research field.

1.6 Relevance to the Curricula of Business Informatics

This thesis explores how the emergence of novelty can be modeled in an evolutionary environment. It therefore contributes to this important question not only in a theoretical way, but combines the evolutionary-economic knowledge with modeling and simulation competences. As a result this combination will lead to a direct application of this knowledge in the form of an agent-based model. The curriculum Business Informatics provides the necessary knowledge base in various courses, the most directly linked ones are listed below:

- 175.036 Evolutionary Economics
- 105.626 Computational Economics
- 101.455 Modeling and Simulation
- 175.035 Information Economics
- 105.622 Computational Social Simulation
- 105.625 Advanced Economics Project
- 105.610 Principles of Macroeconomics
- 105.620 Principles of Microeconomics
- 105.613 Macroeconomics
- 105.628 Econometrics for Business Informatics
- 188.915 Innovation

CHAPTER 2

Innovation and Growth in Economics

2.1 Traditional Growth Models

In the history of *traditional growth theory*, the role that innovation, or more generally speaking, knowledge accumulation plays, is rather ambivalent. It was soon recognized that capital formation can not be the only driving factor for economic growth and that there is also some important technological force. However, how to treat this technological growth, which we would call innovation now, seemed to give traditional growth models a rather hard time. Many models didn't really incorporate it and were forced to handle technological progress as an exogenous process, like we will see very soon in the Solow model as an example.

But before we look into concrete traditional growth models, we will make a slight detour to the general question what technological progress means to economic growth and why it does give neoclassical models such a hard time. A good overview to this topic can be found in 'Innovation and Growth in the Global Economy' by Grossman and Helpman [9] on which this sub-chapter will lean on heavily. One of the most interesting facts about economic growth is, that it exists. And it does exist usually in a very steady way, recent exceptions caused by the current crisis excluded. Nowadays we assume economic growth, take it for granted. We are giving loans based on this assumption, we calculate inflation rates into almost everything and so on. I will for sure not try to start a fundamental debate about the dangers of this assumption and permanent growth in general, but this observation is essential for our starting point. Many countries have sustained a positive and non-declining rate of *per capita income* growth over very long times, even decades. Also the performance in this regard has varied hugely between different countries. It therefore appears that there has to be some 'magical' force in play, that enables to sustain this growth rates. For sure magical is the wrong word, as even most non-economists would suggest that technological progress is the reason for sustained growth and (at least in the western world) high living standards.

It is therefore not very surprising that also classical economists have recognized this circumstance. They proposed that standards of living could not rise infinitely without technological progress that increases the productivity of resources. Also, with diminishing returns to the accumulated production factors there would be no incentive anymore to invest in the long run. “The fact that investment has continued for more than two hundred years since the industrial revolution suggests that technical change has played a major role in the growth process.” ([9]) It is also easy to see - especially based on the different performances among different countries - that technological progress is not a random process, but one that is connected to market forces. So far it is rather easy to agree with these statements for almost everyone - from whichever economic school they might come. But now the question arises how to incorporate this insight into economic models. Traditional growth models had a hard time doing so. The first problem that arose was the assumption, that most technological progress comes from scientific research. But science was pictured to be operating outside of the profit sector of the economy, which would remove industrial innovation from the field of economics. A rather obscure picture, that would even raise the question if economics wouldn’t be completely irrelevant then - especially if you credit to innovation as much importance as I do with this work.

Of course there were critics to this picture, among others also Schumpeter, who we will cover in detail in the next sub-chapter. The main critique was that industrial innovation does not fall from the sky because of more or less random inventions, but is rather driven by capitalistic processes, by the recognition of a problem to be solved that promises profitability. This means that we are talking about intentional innovation, and if we agree that this is the right view, then according to Grossman and Helpman it would be the task of an economic growth model to be able to explain this connection between innovation and growth and also between market conditions and innovation rates. A model with this ability then can for example explain the performance differences between different countries. As a side remark it should be noted that this aim is still far away from finding any explanation how innovation really emerges in the form of a model. But it is a first step.

Solow - Exogenous Technology Progress

So let’s look at concrete traditional neoclassical models or rather their development over time. The first model we look at and which might be even the most famous is the Solow model (1956) with it’s well known Solow residual. The accumulation of knowledge has two main purposes in these models, one could even say is misused a little bit for the purpose of the models.

> First, technological progress may help to explain the ubiquitous Solow residual, the portion of measured growth in national product that cannot be attributed to the accumulation of inputs. Second, technological progress enables capital formation to continue even when the ratio of capital to primary resources begins to grow large. ([9])

The Solow model is for sure a welcome target for all critiques of the neoclassical school. Similar to many other early models Solow was forced to treat technological progress as an exogenous process. The basis of the model is an economy that uses capital and labor to produce one single commodity Z. It can be consumed by households or installed by companies as equipment. With constant returns to scale the productivity of labor is growing exogenously (technological progress). At time t only $1/A(t)$ units of labor are necessary for the same output that at time 0 one full unit produced. The output-input relation is as following

$$Z(t) = F[K(t), A(t)L] \tag{2.1}$$

where $K(t)$ is the capital stock, L is the labor force. $F(\cdot)$ is concave and linearly homogeneous while $A(t)$ increases monotonically over time. The 'intense form' where $z \equiv Z/AL$ is output per effective unit of labor and $k \equiv K/AL$ is capital per effective worker [9] looks then like this:

$$z = f(k) \equiv F(k, 1) \text{ with } f'(k) > 0 \text{ and } f''(k) \leq 0. \tag{2.2}$$

With the simplifying assumption that productivity grows at the constant rate g_A we get

$$A(t) = e^{g_A t}. \tag{2.3}$$

If every household saves a constant fraction s of its income then the aggregate savings are sZ. As the savings are used to finance investment and in this simplified model there is no depreciation, the capital formation equals aggregate savings, so with (2.1) and (2.2) we get

$$\dot{K} = sALf(k). \tag{2.4}$$

To get the development of capital per effective worker we combine (2.3) and (2.4) to a version of Solow's famous equation:

$$\dot{k} = sf(k) - g_A k. \tag{2.5}$$

Figure 2.1 gives a good overview over the movements in k (capital-to-effective-labor ratio). When $sf(k)$ exceeds $g_A k$, k is rising, in the opposite case it is falling. As the curve $sf(k)$ is steeper than the line $g_A k$ close to the origin (marginal product of capital is higher than g_A/s when capital-to-labor ratio is small) and flatter for sufficiently high values of k, the economy approaches a steady state. k will always approach $\bar{k}$, so in the long run the capital-to-effective-labor ratio is constant. Therefore the capital stock grows in the same rate as labor productivity, and also output grows at this rate (see (2.1)). So, everything completely depends on the constant exogenous rate of technological progress g_A. "In this setting neither changes in household behavior nor the introduction of government policy can have any effect on the long-run growth rate." ([9]) If for example the savings rate rises, then capital accumulates faster for a while until a new steady state (with a higher level of per capita income) is reached. The rate of accumulation then settles down again to the one that is defined by the exogenous technological progress.

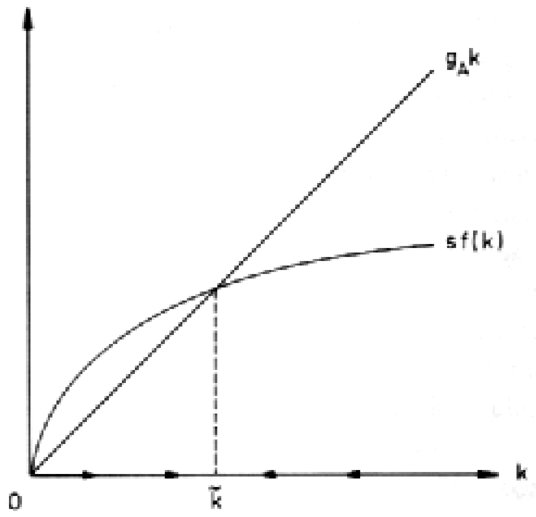

Figure 2.1: Capital per effective worker. Source: Grossman and Helpman [9]

That means that in the Solow model growth can not be sustained without technological progress. Diminishing returns to capital lead to the fact that less savings result from the marginal unit of capital and so on. In the long run as the marginal product of capital approaches zero, the growth of output does as well. Solow himself recognized that there is a case where sustained growth is possible even without technological progress, and that is if the marginal product of capital is limited from below to be a positive number ($f'(k)$ tends to be the constant $b > 0$). This means that returns to capital are then not diminishing but constant.

Arrow - Learning by Doing

However I don't want to go into even more detail regarding Solow. The basic message should be clear, and that is that in the basic setting everything is dependent on the technological progress, which is exogenous to the model. Let's look now on two other classical growth models which try to explain the accumulation of knowledge a little bit differently. We start with Arrow (1962), who proposed that it happens mainly as firms engage in new activities, so it would be a form of learning by doing. What is interesting though is that this learning is happening rather accidental - not as the main purpose but as some sort of byproduct. The state of knowledge at time t is linked to the aggregate amount of investment that had taken place in the whole economy up until this time.

There exist two interpretations of the model. The first one suggests that knowledge is produced by firms automatically when they produce capital goods. This knowledge flows into the public domain (the companies are not able to prevent this) and accelerates the productivity of subsequent manufacturing processes. The model of Arrow modifies the aggregate output function that we already know insofar as that the productivity of labor is now depending on the accumulation of technological knowledge, which is represented by the cumulative investment from the past -

the aggregate stock of capital.

$$Z = F\left[K, A\left(K\right)L\right] \tag{2.6}$$

The Sheshinski interpretation' of this equation would be that the first argument of the input function is the private input of capital by the companies. The second argument represents then the aggregate employment of labor, where we see that the productivity A now depends on K, which represents the state of technology. We already see the big difference to Solow, where the productivity was only a function of time, determined by a constant exogenous growth rate. Here, with the dependence on capital stock (and therefore past actions), it becomes endogenous.

Another interpretation comes from Romer, who views K not as material investment but directly as knowledge. The research process for creating knowledge uses the same inputs in the same ratio than the production of tangible goods. The investment is then in private knowledge (instead of capital) and this is used together with labor to produce the final output plus additional knowledge. The productivity is again dependent on knowledge and obviously this knowledge is coming also in this interpretation from a public stock, to which companies contribute inevitable. Therefore in both interpretations (although in my opinion to a different extend) the overall technological advance is only an unintentional byproduct of companies private investment decisions.

Also independent of the interpretation is the fact that investments by companies will only happen if v_K, the value of a unit of capital (physical or knowledge), exceeds the price of output. Therefore following investment-incentive condition applies (p_z is the per unit cost of capital equipment):

$$v_K \leq p_z \text{ with equalitiy whenever } \dot{K} > 0 \tag{2.7}$$

Therefore the growth of the capital stock looks like the following:

$$\dot{K} = \begin{cases} F\left[K, A\left(K\right)L\right] - \frac{1}{v_K} & \text{for } v_K > \frac{1}{F[K, A(K)L]}\,, \\ 0 & \text{for } v_K \leq \frac{1}{F[K, A(K)L]}\,. \end{cases} \tag{2.8}$$

Without going too much into further details it should be noted that the two different interpretations led also to different outcomes in terms of sustainability of growth. Sheshinksi assumed that returns to cumulative investment in the creation of knowledge would decrease. With a sufficient large K the productivity in relation to capital $(A(K)/K)$ would fall to zero. After the economy has reached a certain amount of capital, there is no reason to invest anymore and therefore also growth ends.

Romer on the other hand came to the conclusion that growth is sustainable if there aren't diminishing returns. Instead of zero, $(A(K)/K)$ could approach a not too small constant. Then there would also exist a long-run constant growth rate. Also the value of the capital stock must approach a constant so that the inter-temporal budget constraint holds. We then have a competitive equilibrium where growth in per capita income is fueled by endogenous technical progress. In this economy government investment incentives can in fact influence the growth rate.

> It should be apparent moreover that such policies can be justified on efficiency grounds. Investment activities create a social benefit that goes beyond the private return that accrues to the investor. This benefit comes in the form of an addition to the economy's knowledge base. Due to the existence of this positive externality from investment, the market allocation entails a suboptimal rate of capital formation. ([9])

Shell - Basic Research

The third model that we are looking at also endogenizes the knowledge 'production', but in a different way than Arrows. As already said, in Arrows knowledge is a product of the investment activities of the private sector, but is thereby created rather accidentally. Shell uses a completely different approach. In his model, knowledge is the intended outcome. But the activities which create it, are very different. Knowledge is now produced in a distinct research sector and the main driving factor is not profit in the first place, but rather "curiosity (and government financing)". This is why Grossman and Helpman call this view 'basic research', as the output here is best to imagine as ideas whose output and potential benefits are rather hard to assign to specific production processes.

Interestingly enough, also the Shell model would end in a stagnation of growth after some time in it's original version. Grossman and Helpman present an alternative version of this model, which is able to sustain growth, like the other 'upgraded' models that we already got to know. This makes it easier to compare this model to the others and also, according to the authors, it highlights what they view to be "(...) the essential insight from Shell, namely that the creation of knowledge often requires the diversion of resources from other activities." ([9])

The production function which we know already very good now looks again slightly different:

$$Z = F\left(K_Z, AL_Z\right). \tag{2.9}$$

The main difference is the subscript Z which indicates that the factors are used in the manufacturing sector. This time the researchers even share their knowledge on free will (compared to the past, where private knowledge was just unpreventable from spilling over to public knowledge) - therefore the productivity increases with the output from the research sector. Interestingly, for the sake of simplicity the production function for knowledge looks almost identical:

$$\dot{A} = F\left(K_A, AL_A\right). \tag{2.10}$$

K_A and L_A are here the inputs of capital and labor in the research sector. It is assumed that this research sector uses these input factors the same way the private sector would do, at the same proportions. Also the private and the research sector are connected and kind of rivaling each other by the constraints $K_A = K - K_Z$ and $L_A = L - L_Z$. Income taxes finance the research expenses and are controlled by the exogenous tax rate t_A. All revenues are used to pay for basic research. These are defined by $T = t_A F\left(K_A, AL_A\right) + t_A F\left(K_Z, AL_Z\right)$, or, as we assume similar factor proportions in both sectors, $T = t_A F(K, AL)$. As $T = \dot{A}$ it follows that

$$\dot{A} = t_A F(K, AL) . \tag{2.11}$$

Capital accumulation in the private sector is defined as the difference between output and consumption demand. According to (2.10) and (2.11) the output $F(K_Z, AL_Z)$ equals $(1 - t_A) F(K, AL)$. This leads to a similar equation for the capital accumulation to the one we already know from the Arrow model, only now depending on the tax rate:

$$\dot{K} = \begin{cases} (1 - t_A) F(K, AL) - \frac{1}{v_K} & \text{for } v_K \geq \frac{1}{(1-t_A)F(K,AL)} , \\ 0 & \text{for } v_K \leq \frac{1}{(1-t_A)F(K,AL)} . \end{cases} \tag{2.12}$$

Again I spare further details, as this whole subsection should only give an overview of the different approaches. Therefore let's immediately look at the implications. The economy will again converge to a balanced growth path. Further, the relative size of both sectors remains the same and the growth of consumption, productivity and capital stock are all happening with the same rate g_A. As also the value of a capital unit falls at this same rate, the capital stock approaches a constant.

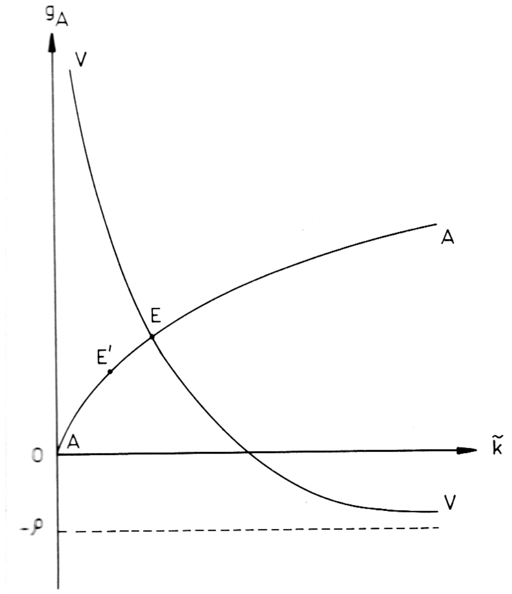

Figure 2.2: Growth rate and capital-to-effective-labor ratio. Source: Grossman and Helpman [9]

Figure 2.2 contains the two curves VV and AA, which show combinations of g_A and k, which is the steady-state ratio of capital to effective labor. Curve VV "slopes downward because the marginal product of capital and hence the incentive to invest grow larger as capital per effective worker becomes smaller." ([9]) Curve AA on the other hand "(...) slopes upward because greater levels of research and thus higher rates of growth can be financed when capital per worker and hence aggregate income are larger. The intersection point E gives the steady-state rate of growth output and per capita income." ([9])

In general, the Shell model has similar implications as the one from Romer. As already mentioned, both models can achieve sustained growth, determined by an endogenous growth rate. Government policy can also here affect the long-run growth rate. If for example the tax rate is increased, the AA curve would shift up because of more budget for the research sector. But on the same time VV shifts down because the incentive to invest got less because of higher taxes. Therefore, growth rate can increase or decrease with a policy like that. If the tax rate was sufficiently small before, the tax receipts will rise. The closer the tax rate is to one, the more likely the receipts will sink with such a tax rate increase. Interestingly enough, the optimal tax rate in terms of technological progress and tax costs will be smaller than the one which maximizes the steady-state growth rate.

Going back to the more general messages that we can take from these models, this is the only one of the three we encountered so far, which clearly depicts the trade-off between technological progress and manufacturing. This is pictured very clear by Shell by having two distinct sectors which share the same input resources. Besides that, we also see that all three models - as different as the basic mindset about technological progress itself, and where it comes from might be - are formally quite similar to each other in the way they are constructed. This is of course also partly because Grossman and Helpman tried to make them in their formalization as comparable as possible. They also all have in common that they need a mechanism which prevents the marginal product of capital from becoming too low. In whatever way this is achieved, it is in all three cases only possible because of technological progress.

2.2 Schumpeter and his Entrepreneur

When we look at the language and the elements used in the previous models, we see typical neoclassical economics; models that are talking about equilibria and constant growth rates. Sure, they are kept very fundamental on purpose, but even with uncountable extensions (and we will come later to more modern successors, which do exactly that) there is inherently no space for the emergence of novelty and for real change; for evolution. Because the language does not allow it (as we will see later in the chapter regarding evolutionary economics). Novelty emerges out of (re)combinations of elements; needs defined connections with their own properties. Even in the models of Arrow and Shell with endogenous technological progress, there is no explanation how capital or labor is transformed into the rise of productivity. Therefore we will turn now to the pioneer in the field of economics, when it comes to innovation: Schumpeter [19] and his entrepreneur.

As a person, Schumpeter is a very interesting case, but I don't want to treat his development regarding his economic and political view in too much detail here. A very good overview of this topic can be found by Hanappi [11]. In his younger years he took sides with the emerging Walrasian faction and its pure market dynamics theory. But already shortly afterwards he would "(...)

depart from his celebration of equilibrating market forces" ([11]) because the historical facts of the real world didn't seem to support them. And facts and describing the real world seemed to be the highest priority for Schumpeter. Therefore it is not surprising, that he turned very fast to the big blind spot of the contemporary "pure economics", and that is the "development problem" as he liked to call it.

Before we go further into the details of his critique, and resulting out of that his own theory and suggestions, there should be stated a slight warning to the reader: Schumpeter has an extraordinary writing style (which for sure doesn't fit everyones taste), but in my view he is more of a story teller. As also Hanappi [11] states, he "never produced a consistent theory. (...) What can be found there is a patchwork of interconnected ideas". Also, mathematical treatments of his theories are missing completely, as he supposedly couldn't handle mathematics too well. But, the reader should not get the wrong idea, that this reduces the value that Schumpeter contributed to the topic treated in this work. Actually, as we will see, his view is a much better foundation for our evolutionary economic model than any of the neoclassical ideas, even though Schumpeter for sure would never have called himself an evolutionary economist.

Critique of Equilibrium-based Economics

As already mentioned, he is criticizing the pure economics (Ricardo among others) especially regarding its lack of development. He says for example, that it "(...) was always its aim, to explain the regular procedure of the economic life based on given circumstances. The development problem for sure wasn't missing completely. But it wasn't vividly perceived and shrugged off rather briefly." ([19], translated by the author from the German edition) He also calls the core of the theory "the statics of the economy". He therefore pioneers a line of critique which will be followed by several other economic schools and streams, most of all the camp of evolutionary economics. We will examine their critical arguments in the next chapter, but Schumpeter actually preempts them in many topics. Especially when it comes to the static world view, which a model of equilibria implicates. He argues that the only thing this static theory defines as "progress" is nothing else than the reaction of the state of equilibrium to the change of data.

> The essential composition of the productive forces is assumed as having somehow emerged and only its variations within the given fundamental forms are described. (...) He (Ricardo, remark of the author) doesn't show us how the swing happened, that has to be happened in technology, the way of economy and general culture, before it comes to the formation and usage of this stock (of production goods, remark of the author), but he immediately begins with the change: If such production goods exist, how does the economic process develop? - He just assumes their existence. ([19], translated by the author)

His critique of course goes into much more detail, but I believe that I have covered the essential parts here. Schumpeter is of course admitting that he is not the first with thoughts in that direction. He mentions James Mill among others who have refined the Ricardian model. According to Schumpeter they are also talking about development, but mainly only about how existing development is affecting the static economy. The question, how this development comes into existence, remains rather untouched. What differentiates Schumpeter from Mill is, according to himself, that he can prove that "(...)the static state does not contain all economic basic phenomenas, but that the life of a stationary national economy differs essentially and in its basic principles." ([19], translated by the author)

He also distances himself from other usages of the word development, such as the "dilletante generalization" of the word or the "unscientific mysticism" that surrounds it sometimes. Interestingly enough he also mentions the perspective which uses development towards the way of thought of Darwin. But, and we will see later in the next chapter that he actually conforms here with most evolutionary economists, he explicitly adds to this statement that it is only problematic if this perspective is used just purely analogous in the economic field. What remains then are according to him two facts: First, that the historical circumstances are changing all the time and these changes are neither describing a circle which is repeating all the time, nor are they pendulum motions around a center. Second, that every historical state is describable out of the previous one. If we cannot explain it in this way then we have an unsolved but not unsolvable problem.

Innovation - New Combinations

To offer criticism should always be only the first part of treating an improvable situation. The second one should consist of providing ideas for this improvement. As already mentioned Schumpeter does not really present a whole consistent and revolutionary framework, which replaces all that was known before. Rather he offers ideas and concepts that are interwoven in some areas and fragmented in others. He starts with asking as "general, as the theory can ever ask: How take these changes place and which economic phenomenas do they trigger?" ([19], translated by the author)

One of the main questions is then, which changes are we even talking about? Not everything that leads to productive improvement is automatically what we would call an innovation. According to Schumpeter the classical theory can handle small changes quite good. The fictive, "intended" but never reached equilibrium even only changes because of changes of data. "If these changes happen outside of social (natural phenomena) or outside of economic but inside social data (war consequences, changes of politics) or if just the flavor of the consumers is changing" ([19], translated by the author), then the typical means of the theory are sufficient enough. "But, these means are failing (...) where the economic life itself changes its own data suddenly." ([19], translated by the author) One example that Schumpeter uses quite often is the construction of the railway, which can be imagined as big economic revolution - nowadays a more striking example would maybe be the introduction of the Internet. In "Business Cycles" Schumpeter even delivers a definition of innovation, and interestingly enough before doing that, he defines economic evolution:

> (...) we immediately realize that innovation is the outstanding fact in the economic history of capitalist society or in what is purely economic in that history, and also that it is largely responsible for most of what we would at first sight attribute to other factors. (...) The changes in the economic process brought about by innovation, together with all their effects, and the response to them by the economic system, we shall designate by the term Economic Evolution. ([20])

Now, let's reveal the probably two most important concepts of Schumpeter regarding innovation besides the entrepreneur who we will treat in the next subsection:

- Innovation is **not driven by the consumer**, but happens in the sphere of the industrial and commercial life. This means that **novelty is created by firms**.
- Innovation is created by the usage of **new combinations** of production means.

The first point seems not particular revolutionary but it has to be stated that the causal structure in the Walrasian school was the exactly opposite. There, every change emerged due to changes in the customers demand. Producers then only reacted to them. Schumpeter on the other hand was always driven by what he perceived as the real economic life, and there it was undeniable that "(...) innovations are usually not taking place in the way that first spontaneously new needs are emerging among the consumers which changes the orientation of the production side (...) but that new needs are instilled into the consumers from the production side, so that the initiative resides in the latter." ([19], translated by the author)

This insight leads us immediately to the second important statement. If the producers are the one who are responsible for innovation, the question is, how are they doing it? Obviously I already gave the answer - new combinations - but we will look now, what that meant for Schumpeter. A more comprehensive treatment of this topic will follow also in the next chapter, as the creation of new combinations is also the foundation for novelty in the model of this work. According to Schumpeter the main definition of innovation is the **setting up of a new production function**. This can be caused by several reasons, namely the creation of a new commodity or commodity quality, the introduction of a new production method, the opening of new markets, a new source for resources or a reorganization of the market, like the creation or destruction of a monopoly. Schumpeter adds: "(...) we may express the same thing by saying that innovation combines factors in a new way, or that it consists in carrying out new combinations." ([20])

As we will see in our treatment of the concept of the entrepreneur it was also very important for Schumpeter to emphasize that innovation and invention are not the same thing but in fact very different, and that innovation not even necessary needs an (recent) invention.

Until now we found little connection points between the traditional models and Schumpeter, and with his narrative style that is anyway a difficult task to do. But the following quote actually provides a great connection which let us look at the traditional model from the innovation perspective, by defining the latter in regards of the monetary system:

> We can define innovation also with reference to money cost. (...) Whenever a given quantity of output costs less to produce than the same or a smaller quantity did cost or would have cost before, we may be sure, if prices of factors have not fallen, that there has been innovation somewhere. ([20])

That means a company can discover innovation by a monetary indicator which doesn't even have to originate of the own industry of this company. The monetary system helps to spill over the benefits of local innovations - just like we already learned from the traditional models.

The Entrepreneur

The entrepreneur is probably the most famous of Schumpeter's concepts and maybe at the same time the most outdated. At least this was thought to be the case until the early 80-ties, when it came back into fashion. And maybe the age of the internet also gives it an unexpected revival, but let's go back to Schumpeter's time first. Schumpeter always brought some ideologic perspective into his theories and the entrepreneur and especially his description in Schumpeter's writing are filled with political motivation. The entrepreneur is almost some kind of heroic figure, the one in a company who makes things reality, who is realizing the innovation. It is of course some kind of celebration of the capitalism in its heroic, anti-feudalist period. Some parts of the description of this hero which we can find especially in Schumpeter [20] are leaving the (supposedly) rational economic sphere at all. They can be peculiar to the reader even when he/she is already used to the story-telling style of Schumpeter, as they are drifting into pseudo-psychological treatments of these individuals.

But all ideologic influence aside there are some interesting characteristics of the entrepreneur. For example Schumpeter leaves the idea behind, that the entrepreneur has to be equal to the owner of a company. Also it doesn't have to be the person who made the technological invention. The entrepreneur is actually not a job position at all, but rather a temporary state, that most people can not claim to be in for a very long time. After realizing the innovation, most entrepreneurs will go back to normal business routine (at least for some time) and therefore lose the claim to be called entrepreneurs. The entrepeneur also is not necessary the one who is financing the innovation. The main characteristic of the entrepreneur is to shift up the production function. When we are looking what this means for growth in general, Hanappi [11] put it the following:

> It was only straight forward that he insisted that only for this core activity a positive interest rate should be assumed. Capital owners get a part of the interest rate which entrepreneurs earn, and if the latter do not innovate, then the interest rate is zero. ([11])

But, after World War 2 Schumpeter himself started to struggle to connect his former (at this time probably rather fitting) concept with the new reality, which consisted of big international companies which steadily grew in size and power. Anyway it was hard for him to find political ground as socialism and Keynsian economics where on the rise. Only with the conservative rollback in the early 80-ties his entrepreneur came into fashion again. But of course there was mainly just a lot of political misuse of the term, such as “entrepeneurial spirit”. To the actual economic and political situation, his theory could not provide much anymore. Actually even the current situation with first the rise and now the omnipresence of the Internet, would be more fitting to his concept than the time back then. Right now, the character of the entrepreneur in fact can be found again, with a lot of successful startups which started extremely small and just with a good idea and are global players now. But still, even if we agree on the existence of the entrepreneur in the current time, this does not really help us in finding a consistent theory that covers the whole economy.

There are other parts of Schumpeter’s theory which seem to be of more relevance even in current ages. Mainly what he has to say about business cycles is of high interest - not for nothing one of his most important books is titled after them. The rise and fall of companies is according to him highly related to innovations. When we put the entrepreneur aside we nowadays still find a lot of descriptions of the life cycle of a company that sound rather similar, even though maybe not with such a strong emphasis on innovation. Schumpeter suggests to set firms equivalent to their innovations. This seems to be a rather logical way of thinking, as companies which are not able to present new products are supposed to lose the race and therefore collapse. But his idea is even more radical: He wants that we suppose that every innovation is embodied in a new firm founded for exactly this purpose. Interestingly enough, in the age of the Internet this might be more true than it had been over the centuries before.

Even if one rejects this assumption, the consequences that Schumpeter draws seem hard to argue about: Firms do not exist forever. They are part of a life cycle in which they pass different stages. According to Schumpeter the life goes out of them when their purpose has been fulfilled, which means when the innovation is not new anymore. Even if we argue that a company is able to produce more than one innovation, this is mostly semantics. It does not change the statement that old and inflexible firms are often having a hard time to keep up the pace in innovating compared to fresh and flexible ones. This picture couldn’t fit more perfectly to the views of evolutionary economics. One has to be very careful with direct analogies from the biologic field (as we will treat in more detail in the next chapter) and especially the struggle of companies introduced to a lot of misuse of “the survival of the fittest”. Nevertheless, for someone who wants to model in an evolutionary economic perspective: The image of a heterogeneous pool of companies, each in a different stage of the business cycle (or their life cycle), struggling for their very existence and the only way they can survive is by innovating... - As I said, one has to be careful, but it is indeed a nice foundation for an evolutionary economic model. With some adjustments also for the model of this work.

2.3 Modern Mainstream Innovation Models

Of course neoclassical growth models have been advanced since the foundations that we discussed earlier. I neither have the space nor the contentual focus here in this work on this topic to show all the possibilities that a real, bloated neoclassical model would incorporate. But I will lean instead on Aghion and Howitt [1] and show how a basic model setup in a 'Schumpeterian' approach of the mainstream could look like and point out the most interesting possible extensions regarding innovation which come at least in the origin of their idea close to what we want to do in evolutionary economics. Formulas I will try to reduce to a minimum, because once started, usually you have to go all the way, as otherwise the single fragments don't make too much sense.

The Basic Model

In this basic model, growth is generated by random improvements which are a result of uncertain research activities. These innovations are vertical and improve the quality of the product so much, that the previous technologies or products are made obsolete by it. Aghion and Howitt call this *creative destruction*, which has positive and normative consequences. Positive insofar, as there is a negative relationship between current and future research, and therefore has a unique equilibrium of balanced growth. But it can also result in cyclical growth paths. On the normative side, current innovations have negative externalities on the prevailing producers. This negative relation is called *business stealing effect* and can lead to excessive growth, something that we didn't see in the classical models.

$$z = Ax^{\alpha} \tag{2.13}$$

The first big difference is that this basic model abstracts completely from capital accumulation. We have an output function (2.13) that is depending on the input of an intermediate good x $(0 < \alpha < 1)$ and an technology parameter A which is raised by the constant factor $y > 1$ by inventions of new varieties of the intermediate good, which replace the old ones. The fixed stock of labor L in the society can either be used for producing the intermediate goods or for research activities n, so these two possibilities are obviously in competition to each other:

$$L = x + n \tag{2.14}$$

The 'arrival' of innovations is modeled by a Poisson process and therefore the Poisson arrival rate is λn, with $\lambda > 0$ as the productivity factor of the research technology. The basic principle is that in the race for innovation, the company which succeeds first can monopolize the intermediate good until it is replaced by the next innovator. There are two ways of positive spillover:

First, the monopoly rents are less than the consumer surplus of the intermediate good. Second, only the invention makes it possible for the other companies to start research on the next innovation. Obviously there is also the negative effect of business stealing, as the new monopolist eliminates all possible surplus of the previous innovation, as this is obsolete then.

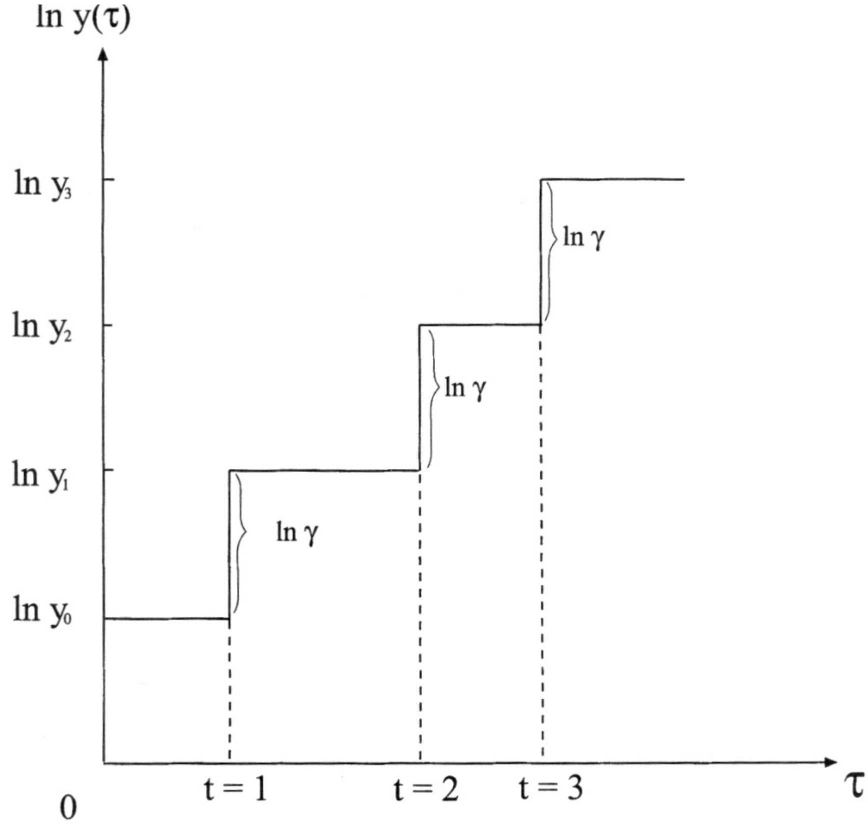

Figure 2.3: Time path of the log of final output - A random step function. Source: Aghion and Howitt [1]

As already mentioned, the model has an unique steady-state equilibrium where the division of labor between research and manufacturing remains unchanged. Then, growth is stochastic (because of the stochastic arrival rate) but balanced (see figure 2.3). The average growth rate can be either too low or too high to maximize the society's welfare. This depends if the negative or the positive externalities that we already mentioned before are prevailing.

It is also possible that there are unbalanced equilibrium growth paths, where the intensity of research switches every period between a low and a high value because of the anticipated amount of research after the next innovation. If the companies expect that this amount is low, they will invest a lot into current research, as it is to be expected that afterwards they have the monopoly for a rather long time. The opposite is true if the expectation is of high research and therefore only a short time of supremacy. We then get some kind of oscillation between these two values;

a process that can be seen as the first basis for incorporating business cycles. Also this opens the theoretical door for a *no-growth trap* where everybody expects so much research that nobody is researching anymore.

Critique of the Basic Model

Unlike many of their predecessors, the authors are well aware that this basic model lacks several important features. According to them, it shares the a number of limitations with other endogenous growth models. Among them, one of the biggest is for sure the reliance on steady-state construction. An observation that evolutionary economists would certainly share. The main problem is that certain assumptions have to be made, to make this steady state possible. And these assumptions "(...) are quite severe and have nothing to recommend them except for tractability." ([1]) Further "These strong assumptions rule out important phenomena, and answer important questions by mere assertion." ([1]) A circumstance that could be often said about neoclassical models in general. But in the concrete case the authors especially highlight that there is no foundational change in the construction of the economy possible. We know for example from our own history that resources get more and more reallocated first from the primary to the secondary sector and later to the tertiary one. Such a change would not be covered in the model, it just assumes that technological growth is to be equatable with scaling up the economy of what it was years ago. But in reality, when we experience technological paradigm shifts, we can have long phases of rises, declines or fluctuations almost equally likely in between.

The second big limitation (among others, that we can't cover here) and probably even the biggest problem from an evolutionary point of view, that the authors acknowledge, is the description of knowledge as parameter A. Technological knowledge is here treated like any other aggregated production factor. But, as Aghion and Howitt stress correctly, new ideas (or combinations) are what is driving technological growth. It might make sense for labor or capital, but how can you aggregate ideas? It seems more appropriate to use some models of collective learning, where the agents for example learn about the true parameters or function describing technology by experience, experimentation, intuition or creative extrapolation and it is not sure that these individual learning processes converge to the truth.

Extensions of the Basic Model

As we covered now some points of critiques it is about time to see what the authors have up their sleeves to extend the model and erase the shortcomings. First we will look at what they call "immediate extensions" to enhance this basic model without changing it completely. Afterwards we will briefly look at some broader adaptations. I will focus only on a few of them, mainly these which are related to innovation and novelty, which is not necessary related to the general importance of these extension for the applicability of the model itself.

The first one would be technology transfers and cross-country convergence. In the basic model, two countries will always diverge regarding their development of the log of GDP. As we have a random walk with drift, even with exactly the same parameters we would most likely see divergent development paths. But this is not exceedingly realistic. We can extend the model so that there exists convergence, which means that not so developed countries tend to grow faster and catch up. Countries with a knowledge level A below the average will achieve bigger innovations than the average and for countries with a higher A it is the other way round.

Another interesting variation is to introduce nondrastic innovations. Until now we assumed in the basic model that innovations are drastic; that means that the previous monopolists and their innovations do not stand in competition to the current one, because the new generation is so much better. In the nondrastic case, the previous incumbent could make a positive profit when the current one is charging the monopolistic price. Therefore he will set "(...) the maximum price that gives the previous incumbent non-positive profits and satisfies all the demand at that price, leaving none to the previous incumbent. The previous incumbent could make a positive profit if and only if a competitive producer of consumption goods could produce at a smaller cost using the previous incumbent's intermediate good, buying the latter at a price equal to its average cost of production w_t." ([1]) The comparative statics results derived from the drastic case are valid also in this case, therefore so far not much changes. It should also be noted that with the introduction of contracts between the current and previous incumbents higher monopoly profits could be reached, for example the previous one could sell his patents to the current one and therefore not compete.

It's getting more insightful when we also introduce an endogenous size of innovations. That means, that companies can decide how big, or more precisely how novel their innovations are. Of course this comes at a cost: The bigger the innovation, the harder it is to discover. In terms of our model that means that the arrival rate is inversely proportional to the degree of novelty of the innovation. It is not a big surprise that under the assumption of drastic innovations are too small than would be optimal. The main target is to steal the business, and being the monopolist is everything that matters. In the nondrastic case, things are looking a little bit different. The new incumbents have to 'distance' their product from the predecessor and therefore the innovations have to be bigger. Otherwise that would reduce their profit margin drastically. Nevertheless this does not necessarily mean that the size of innovations is optimal in this case, it might still be too small.

Innovation and Business Cycles

To make the model applicable for mainstream economics, there is no way around incorporating capital into it. But before the authors are doing that, they introduce another extension, which might be considered more as a generalization. We are talking about multisectors. There still will be only one final product, but different intermediate goods. And for every intermediate good we can create a market like we already know it from the basic model. However what we now get is *technological spillover*: Innovations in one sector enhance the technological knowledge in the other sectors, possibly increasing the productivity there even more than in the original sector. Also the growth process will change drastically: Instead of the discontinuous jumps we got to know in the basic model, the aggregate output will now grow much smoother, as there is now a continuous flow of innovations in the whole economy. “(...) The law of large numbers prevents the uncertainty of the innovation process within each sector from showing up at the macroeconomic level. Thus in a steady state growth is constant at each point in time rather being stochastic.” ([1])

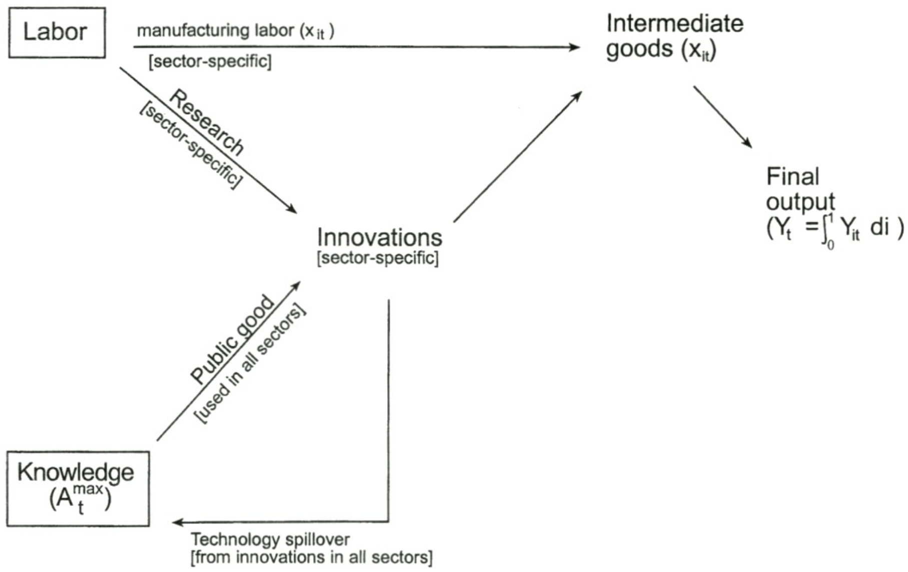

Figure 2.4: Schematic representation of the multisector model. Source: Aghion and Howitt [1]

Figure 2.4 shows a schematic representation of the economic activities in the multisector model. Every sector has it’s own research activities, and the companies within that sector still try to claim the monopoly for the intermediate good that is produced in that sector. The arrival rates in the different sectors are independent of each other, but the innovations all draw from the same pool of shared ‘leading-edge’ technology.

As announced before, it is now time to introduce capital in the model. In reality, innovations are usually represented by a durable good, which means either physical or human capital. Improvements of technological knowledge often transfers to improved machines and equipment. Also the research itself is obviously dependent on capital. Physical capital (computers, instruments, and so on) is needed as well as human capital, where the latter is obviously especially important when it comes to scientists and engineers. The model gets changed insofar as that not labor, but capital is now the input factor for the intermediate goods. Also the final output can be used not only for consumption anymore, but also as an input for research and to produce capital goods.

> The main result of this chapter is that capital accumulation and innovation are both essential inputs to long-run growth. More innovation stimulates capital accumulation by raising the marginal product of capital. More capital accumulation stimu-

> lates innovation by raising the profits accruing to a successful innovator. This result runs counter to the conventional belief to the effect that innovation alone determines the long-run growth *rate* while capital accumulation determines only the *level* of the long-run growth rate. The chapter shows that the problem with the conventional view is that it ignores the role of capital as an input to research. ([1])

In our further discussion regarding evolutionary economics (see chapter 3), we will come to the question if there is a certain degree of novelty which we can attribute to innovations. Also we will see, that there is some difference between groundbreaking *basic innovations* and *follow-up innovations* which are just refining or reusing the basic innovation (or strong and weak emergence, however you want to call it). A similar distinction we also find in Aghion and Howitt [1]. Here they call it *fundamental* and *secondary research*. "Some of them are more secondary than others, in that they do more to bring about the realization of possibilities that have been created by previous innovations, and less to open further windows of opportunity. In reality, every innovation is fundamental to some extent and secondary to some extent." ([1])

Nevertheless the distinction is obviously important. Therefore the model will keep it simple and assume that there are only these two different kind of innovations. Also another concept is introduced: In reality not all discoveries are happening only via research. Especially secondary innovation are often triggered by problems or discoveries which are happening during the production process. This is called *learning by doing*. Generalizing, we can say that fundamental research creates new product lines or windows of opportunity, while learning by doing is improving them or filling them up.

The conclusions of a model which is incorporating these ideas (see Aghion and Howitt [1] for details of the model) are the following. First, in the case that learning by doing is purely external to companies, it could happen that too many resources are used for research instead of learning by doing, which is hurting the growth. In the contrary case that quality improvements are solely based on internal learning by doing, then research always has a positive influence on the growth, as research is more forward-looking than learning by doing. Another topic which pops up in that connection is worker mobility. Growth can be increased when workers can switch faster to new product lines or to research. Although this is called the *Lucas effect*, the authors disagree with his view that growth increases because the higher mobility leads to a higher pace of learning by doing. According to them, first and foremost it increases the profitability of research, as having more workers starting faster in the new product line is beneficial for the innovator.

As Aghion and Howitt are striving to incorporate Schumpeterian views into their models, there is also an attempt to include his famous business cycles. In the economics literature, trend growth and cycles and fluctuations were treated completely separately, and models of (exogenous) growth ignored business cycles completely. This is also highly connected to underlying basic assumptions: Neoclassical growth models assumed perfectly working markets, which cycles literature obviously does not. Also the *neutrality of money* is an important aspect: With endogenous growth models the assumption of money neutrality, which was a cornerstone for neoclassical economics, does not hold anymore. Quite the opposite is true, as we have strong nonneutrality: “(...) the real impact on aggregate output of a temporary monetary shock *increases* over time as a result of learning by doing.” ([1])

The authors analyzed the relationship between growth and cycles from both sides. For example a temporary boom (a cycle) increases temporary the output and research and therefore has also permanent consequences: The economy will be lifted to a higher growth path. But research might not rise only in a boom: There are a lot of theories suggesting, that the slumps are actually having an even bigger positive influence. Several reasons come to mind: Inefficient companies are ‘filtered’ out, making the employees free for more efficient (and therefore more researching) companies. Or maybe there is just less need to produce (no demand) and therefore more capacity for research activities. Higher efficiency is needed to survive, so research might be the only answer to avoid the collapse. As reasonable as it sounds, the authors are warning to see too much positive in slumps. Evidence to support these ideas is according to them rather thin. Their suggestion is that the topic can’t be analyzed so isolated. In general, in times of recessions, there are not only slumps, it’s just a big chaotic mess of ups and downs. And large swings mean large uncertainty and this in turn means delaying or stopping of investment and research.

Of course also the other direction is interesting: How can trend growth influence the business cycles. Here, the authors introduce the concept of *GPTs*, *general purpose technologies*. The idea is that groundbreaking innovations like the steam engine or the computer lead to a long time of adaption, even lasting for centuries. Their arrival is a big shock to the economy and they lead to a wave of further innovations. We will stumble upon that concept again in the next chapters. These technologies can in the beginning lead to stagnation or even slumps. One reason is, that their full potential first has to be figured out. Often this is not possible without certain follow-up innovations. Another reason according to them is *social learning*. Firms are learning from each other, and therefore their speed of learning depends on how easy it is to observe the competition. When such a GPT arrives, we see a highly complex web of connections between these companies, which struggle to do their own costly experimentations with the new technology, trying to get the best benefit out of it and on the same time peeking at what the competition is doing. Droughts and snowball effects will vary in this time and therefore create fluctuation (cycles).

But the little empirical evidence (not many innovations are so big to be interpreted as GPTs) does not really support the theory that actual slumps are created. New GPTs could lead to measurement problems and therefore only give the illusion of a slump while in reality the growth is perfectly fine. One of these problems is to only measure the production of consumption good as an indicator of growth, while production of knowledge also contributes to growth.

CHAPTER 3

Evolutionary Economics and the Emergence of Novelty

3.1 Evolutionary Economics in a Nutshell

The title of this section is, deliberately, a little bit misleading. Because, as we soon will see, to describe evolutionary economics ‘in a nutshell’ is not only a very hard task, it’s completely impossible in my opinion. The picture of what evolutionary economics is about, is very fuzzy and diverse, even among its advocates. For those who work in this field, this is no problem at all, as it perfectly fits the own paradigm: Different ideas, models and theories evolve and are in competition. Science is in constant change, as is the world it wants to describe. This is nowhere more true than in the social context, and especially in economics. Nevertheless our standard paradigm is still about equilibrium everywhere, accompanied by the equally unrealistic assumption of rational expectations (which are at least according to Potts [18] only symptoms of a deeper ontological problem, which is the concept of the field - but we will come back to this later again). For the sake of simplicity I will summarize these concepts as ’standard economics’ and will only deviate from this term if an author explicitly highlights some specific characteristics of this paradigm with a different term.

For outsiders on the other hand, it is almost impossible to grasp what evolutionary economics is about, and what its foundations are. Even for someone who has read Darwin [4] and Dawkins [5] and therefore claims to be quite familiar with the biological meaning and implications of *evolution*, it can be hard to find the connections between these biological concepts and a scientific framework that could be used in economics. Part of the reason is definitely the problem that ‘evolution’ and ’evolutionary’ are used far too often too randomly as empty words. Also, they are too often used for direct analogies, which compare apples with windmills, like companies that fight for the survival of the fittest, and so on. The papers (and the ideas they contain) that we will now discuss, can help a lot in drawing an own mental picture of what evolutionary economics is really about.

After reading several of these statements which propagate the idea of evolutionary economics in quite different ways, three main approaches and/or questions crystallize out, with the first one being the most elementary of course:

1. What is evolutionary economics and how differentiates and delimits it itself from standard economics? And implicitly, as we are talking about advocates of that new paradigm: Why and how is evolutionary economics better?
2. What role should *Darwinian* concepts play in an evolutionary economic theory? Should they be directly applied, only as analogies or is there even a general ontological basis, something like *Universal Darwinism*?
3. Which new concepts (often already implicitly used, but not explicitly stated yet) result from using evolutionary economics?

What is Evolutionary Economics?

In some way, all of the authors deal with the first question, mostly beginning with a critique of the standard economics. However I want to point out Arthur [2], Dopfer [6] and Potts [18]. Arthur interestingly enough never speaks explicitly about evolutionary economics, but rather about out-of-equilibrium economics. He clearly comes from the side of agent-based modeling, but is in my opinion right in saying that all the economic branches like complexity economics, agent-based modeling or adaptive economics are nothing more than different manifestations of one new economic path, only with different specialties and nuances. We would call that path evolutionary economics; he prefers out-of-equilibrium economics, simply to point out the differences to equilibrium economics, which he heavily criticizes. Nevertheless he states that "evolution emerges naturally from the very construction of such modeling" ([2]). The key element of his proposed approach is *perpetual novelty*. Here he states, not surprisingly, that standard economics cannot produce perpetual novelty with its equilibria, but an evolutionary approach can.

However it gets more interesting when he points out two more concrete problems of standard economics and how they could be solved with the new approach. The first one is the problem of *multiple equilibria*. Standard economics can be used to identify different consistent patterns, but cannot tell us how one is chosen – there is an indeterminacy. Arthur proposes therefore to see the problem not as the selection of one of several god-given equilibria, but as one of constant equilibrium formation. It then becomes a dynamic process with random events, where different equilibria can emerge under the same conditions. The second problem is similar, it is about expectational indeterminacy. To put it short, Arthur describes his famous *El Farol bar* problem, where rational expectations wouldn't work at all (the outcome they predict could never be realized). But with an agent-based formation of expectations (internal model building of the agents), this problem can be solved. In summary, Arthur sees this out-of-equilibrium theory not in competition to equilibrium economics, but rather the latter as a special case of the first. This is of course a clever way to put the new theory ahead of the old one.

Dopfer on the other hand uses a more epistemological, but to me even more enlightening approach to describe the foundations of evolutionary economics. He speaks in terms of paradigms (unquestioned believes that determine the way of building theories) and compares the standard economic paradigm (he calls it a mechanistic, neo-classic economic paradigm) to the new evolutionary one. He compares the ontological assumptions of these paradigms and describes the evolutionary one as *histonomic*, in contrast to the *nomological* program. Histonomic means that there is a local repeatability of phenomena, but there is also a historical context, which makes global repeatability impossible (at some point, there is a change). He then describes three different mechanistic axioms and compares them with their evolutionary pendant.

The first axiom says that every matter-energy actualizes itself, and it is doing this in one specific way. The evolutionary paradigm agrees on the first part, but doesn't believe in one specific way, which could be called some universal law. Instead, not only are there multiple ways for actualization for a given point in time, but the set of those different ways also changes for every point in time. In a mechanistic paradigm all single particles of matter-energy follow the same path (trajectory). As the law of this trajectory is invariant, for a given starting condition, there is no variety in the actualization. Following the thought of Epikur and his *clinamen*, some of the particles "at times quite uncertain and at uncertain places, (...) swerve a little form their course, just so much, that you might call it a change of direction". The second axiom says that all particles are isolated. This is based on Newtons mechanics, where every particle follows its own laws. There is no feedback between micro- and macrobehaviour. Dopfer doesn't really declare the evolutionary counterpart to this, but it's obviously the complete opposite, so that there is feedback, the consequences of the actions of particle recoil to itself. The third axiom then states that particles are continuous. The classical model describes only continuity, not discontinuity of movements. In Newtons mechanics, there is no causality, no reason given for the beginning and the ending of a movement. Of the three phases "origination", "continuation" and "termination", only the middle one is described endogenous, the other two are exogenous. I guess it is no big surprise that in the evolutionary paradigm, all three are endogenous.

A somewhat similar, because also ontological attempt to unify the view on evolutionary economics comes from Potts. In his book "The New Evolutionary Microeconomics" Potts not only criticizes standard economics, or, how he calls it, *neowalrasian economic theory*. In the course of his critique he also tries to unify the perspective of all 'heterodox schools', by saying that they in fact all criticize the same main issue of the 'orthodox school', but are either not really aware of it, or just come from so different starting points, that they look at the same topic from too different angles. This is a rather brave attempt, as Potts includes in these heterodox schools among others the Evolutionary, Realist, Behavioral, Institutional and even Austrian and Post-Keynesian school. So, even as it is obvious that all these schools criticize certain aspects of standard economics, it is at first glance hard to believe that all their critique can be unified into one fusion. So what is this mystical key aspect underlying the neowalrasian economic theory that is the target of this unified critique? According to Potts it is the concept of the field.

The neowalrasian economic theory is built on a real field R^n. There hasn't been much critique to this concept, as it was very successful in mathematics (foundation of integral and differential calculus), and therefore it is also the underlying principle of modern science - natural sciences especially. This tremendous success made obviously a huge impression on economists, who ever since tried to make economics more 'scientific' and adopted this concept. But what is the problem then and why shouldn't the field concept be as successful in economics as it is in physics? Potts argues that the main problem with this concept is, that it does not recognize connections as own elements. Rather in a field it is automatically assumed, that all elements are connected with each other, or as we can see in general equilibirum theory, "that everything is a function of everything else must be understood to be literally true" ([18]). For Potts, in physics this is no problem, as in a gravitational field indeed everything affects everything. But for economics it seems rather unrealistic to assume this universality. As soon as it can be shown that not everything is connected with everything, there exists no legitimation for the concept of a field anymore. Rather should the focus lie on specific connections, which become elements themselves and can have their own properties, because otherwise

> Interactions, knowledge and structure are specific connections between points in space and therefore the very existence of these concepts is excluded by the assumption that all points relate, a priori, to all other points directly; that is, with a single mathematical operation... In a field, 'interaction' does not mean anything because there is no example of non-interaction - everything interacts with everything else. ([18])

But what is then the answer to this problem? How can you give these connections meaning? According to Potts there is only one way to realize this, and that is by the usage of graph theory and complex systems. Therefore he suggests that there are only two hardcore propositions which build the ontological backbone of evolutionary economics:

- Evolutionary-HC1: There exists a set of elements.
- Evolutionary-HC2: There exists a set of connections.

([18])

While this seems rather trivial and the first proposition doesn't differ from what we already know from standard economics, the implications of the second one should not be underestimated. It sets the connections on the same level of importance as the elements. We will see that this is especially important regarding innovation, the topic of this work. Because combination and recombination of elements is strongly dependent on the connections of these elements. We can observe in reality that novelty is only very seldom created by the introduction of completely new elements. It's most of the time even not about **which** elements are combined, but rather **how** they are combined - for example in which order, with which manufacturing process, and so on. We will go into this topic much deeper in the next sub-chapters, but this anticipation should only highlight the importance of Potts' statement for this work.

The Role of Darwinism

I had to go into a little bit more detail in the previous three explanations, because in my view they are the core explanation what evolutionary economics is founded on: Non-equilibrium, agent-based internal model building, a strong focus on the connections between the elements, a histonomic paradigm and possibilities within this paradigm to allow deviation, variety and change and additional give the opportunity to even explain those phenomena from within the model. But of course, this is still a very general view. We all know that our modern view on evolution originates very much from Darwin's revolutionary theory [4], so the question arises: To which extend can we apply the concepts of this theory to an economic framework? Witt [23] gives a very good overview (one could even say hierarchy) of the different possibilities – "with or without sharing the Darwinian ontological assumptions".

The first approach is the most direct one: To apply the theory of natural selection directly to human economic behavior. The argumentation of this approach is that human behavior itself is a product of biological natural selection, because human undergo these natural selection and so there has to be some 'fitness' in certain behavior. This approach is easy to criticize, because first of all, this could only be true to genetically derived behavior. Secondly, this might have been true for early human civilizations, but nowadays there is no such biological pressure anymore. Also most of the behavior in economic terms already changes within one generation. If anything at all, this is a cultural, and not a biological evolution.

The second approach is the one I have already criticized earlier: It doesn't claim a common ontological basis but uses evolution rather for analogies or metaphors. This can be very problematic if this heuristic usage of the Darwinian concepts is done on self-purpose and doesn't seem to be justified by an adequate interpretation of the economic phenomena, which is the case most of the time in my view.

With the third approach it gets more interesting. This is especially proclaimed by Hodgson [15], who kind of borrows the term "Universal Darwinism" of Dawkins. The subtitle of his paper is "from analogy to ontology" and that is exactly what he wants. In contrast to the previous approach, there is the strong belief that all open systems work on the basis of Darwinian concepts. According to Hodgson "an adequate evolutionary economics must be Darwinian, at least in the fundamental senses" ([15]). He doesn't want to 'borrow' Darwinism from biology ("biological imperialism"), because for him Darwinism is a universal metatheory, the biological application is only one of many. As an example he mentions that instead of genetic replicators, there could exist other replicators in other fields, like human communication, conformation and imitation in a social context. Witt criticizes Hodgson's view, but not because of the ontological assertion (which he shares), but he questions the general validity of Darwinian principles for all evolutionary phenomena.

From my point of view they mean both pretty much the same. Hodgson sticks more to the terminology of Darwinism, but as Witt points out correctly, he in fact uses them for describing the principles of *variation, selection and retention*, which are already abstract reductions of the neo-Darwinian theory in biology. So if Hodgson would stop to call his approach 'Darwinian', content wise they both would be on the same page. Witt himself calls his 'fourth' approach the *continuity hypothesis*. The main question is according to him, if the different domain-specific forms of evolution have anything in common, some generic features, which can be observed wherever evolution takes place. Witt thinks, that he has found this common ground in the two generic features of evolution: the emergence and the dissemination of novelty. We will obviously come back to this assessment in the next sub-chapter, as novelty is the key topic of this work.

New Concepts

As we can see, although I have now tackled the question of 'What is evolutionary economics, and what is it not?' from quite different angles, there can't be given some clever one-sentence answer. However by studying different perspectives and opinions, one can create a picture in ones mind, that may be a little bit fuzzy on some spots, but in general gives a good overview. The same is true for the last question that I have raised at the beginning, namely, what new concepts could result from this new paradigm. I will give only a very brief outlook by mentioning three interesting ideas. One is from Hanappi [10] who suggests 'evolutionary economic programs' as antithesis to paradigms. According to him they are necessary as a mean to synthesize the world of formal language and the world outside language, like computer programs are written in a language, but perform something outside this language. Dopfer, Foster, and Potts [7] go in a completely other direction, but maybe even further. They want to enhance the micro-macro architecture with a *meso*-layer. According to them, from an evolutionary perspective, the macro layer is not simply the aggregation of the micro units. Instead, the macro consists of the population structure of meso, and micro consists of the individual carriers of rules. But what is meso? The authors argue that an economic system is a system of rules. And a meso unit is exactly that: a generic rule and its population of actualizations. They believe that evolutionary economics need the meso perspective, because it "deals with system dynamics head on in terms of structural change and open system process; micro-macro does not" ([7]).

Potts himself, whose book was already mentioned before [18], also has some own ideas about what consequences arise from an evolutionary economic perspective. In his case they are a logical consequence of what in his view is the difference between 'standard economics' and evolutionary economics. As according to him the distinguishing feature of evolutionary economics is the concept of connection, there is only one answer how to realize this: graph theory and complex systems. His first two propositions, namely that in the evolutionary economics there exist elements and connections as equally important parts, were already mentioned before. But there is one more:

- Evolutionary-HC3: There exist systems. Systems are the basic objects of theory and units of analysis.

([18])

According to Potts everything is a system - be it a company, a product or whatever. Systems exist of other systems and are thereby always complex systems. In comparison to the static view of a field, in the evolutionary picture (potential) change can be represented very easy, by moving within the state space of the system (see figure 3.1). For example for a specific state of elements

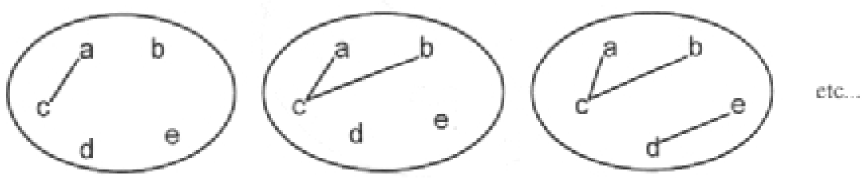

Figure 3.1: State spaces of a system. Source: Potts [18]

and connections, every other state that can be reached within one step can be defined easily as the neighborhood of this state by removing or adding exactly one connection. This view shows that particular states are easier to reach than others and therefore more likely. In a field this is not clear, because without the concept of adjacency every point in the field can be reached by a single change and therefore leaves no room for search strategies or adaption. Evolutionary economics on the other hand sees the incompleteness of the spatial and temporal dimension as a key characteristic of economics in reality.

> These dimensions do not exist a priori, but are created in the process of economic coordination. By looking forward into the future, economic agents thereby create that future. ([18])

We will see whom of them is right, but I'm excited how evolutionary economics will develop in terms of its ontological aspiration. At least one thing is for sure: New ideas will evolve.

3.2 The Importance of the Emergence of Novelty in an Evolutionary Context

Whether we consider economic or biological evolution, the emergence of novelty plays an undeniable big role - many would even consider it the key feature of an evolutionary theory. In this section I will try to examine why exactly this is the case - Why is the emergence of novelty actually so important? We will answer the question at the same time in a general evolutionary context (coming from the famous Darwinian theory in biology) and of course also with a special focus on economic evolution. We also have to ask how, as soon as one accepts the importance of the economic emergence, there still can be a mainstream economic theory without any evolutionary aspect? Of course we will not discuss the (historical grown) reasons for the ongoing prevalence of equilibrium-based models despite their clear failure in the current crisis. But we can at least present some arguments why the emergence of novelty should be important to any economic theory and concluding from that, why every of these theories should be built on the basic principle of evolution.

Novelty in Biology - In the Footsteps of Darwin

Let's start with a very quick side-glance to the field of biology. Even though he didn't know anything about genes back then, Darwin [4] and his famous "Origin of Species" was the foundation of evolutionary theory as we know it. It is common knowledge that there the interaction of changing environmental conditions and 'random' *genetic mutation* which fits these changes better than the 'normal' genetic code of this species is leading to the emergence of genetic novelty, which means new species in the long run. Before we go one step back and look at a more general view of what novelty exactly is and how it can be pictured in general and economic terms, first of course the question rises how far the similarities between biological and economic concepts can go. We already had this question in terms of evolution in general, but it is also interesting to look just at the emergence of novelty. In the field of biology one of the main arguments of critiques or disbelievers of the evolutionary theory was always that it is impossible that complex systems like the eye can arise out of nothing. Also their line of argumentation is, that intermediary steps like 'half of an eye' don't make sense.

Evolutionary biologists like Dawkins [5] usually bring up the counter-argument that no matter how big the transformation and how complex the resulting system might be, it is always possible to reach it with enough small steps (genetic mutations). And of course intermediate steps can give an individual already enough advantage that every further mutation in that direction will lead to more superiority. Actually the eye is a perfect example because if we think about an organ, which is not as developed as our eye but allows its owner to already see a little bit, however blurry that might be, then this individual has a huge advantage over all others who don't see at all. So, while there is no reason to doubt evolution in the biological field these almost 'infinite' small steps raise some more questions. In terms of biological evolution is mainly then: 'Where does one species start and where does the other end?'. Because as we already said, there are no big sudden genetic jumps from one generation to the next one.

With a small glance back to our economic starting point this implicitly seems to raise problems. Are there really almost infinite small steps between the old technology and the new one? Or aren't innovations exactly these big jumps which are so unimaginable in the biological area? Even though in my opinion 'Universal Darwinism' [15] is only acceptable if you reduce it to the basic evolutionary principles (as Witt [23] suggested) and too direct analogies are always dangerous, one shouldn't underestimate the similarities. Also in the field of biology there are both, calm and very revolutionary phases. In the latter in a rather short span of time (for evolutionary biological standards) a lot of changes are happening and therefore new species are emerging. In hindsight to evolutionary economics Hanappi calls the calm phases *regimes* (we will come to this picture later again in more detail regarding economics) and the gaps between these regimes revolutions:

> A species not only originates in the stabilizing feedbacks that the environment produces in a population (basically a stable growth process, call it 'Fast Process' of a certain stabilizing 'regime'). At the beginning of a regime new species emerge as the result of a much shorter process, their origin (...) thus is to be found in the 'revolutionary dynamics' following a far-reaching extinction of other species or natural catastrophes. An often cited example is the Cambrian revolution, which clearly has led to an enormous variety of species, which only in a later stabilizing phase was reduced by Darwinian selection processes to fewer species. ([13])

We know this 'snowball effect' only too good from the history of inventions and innovations. Every revolutionary new technology has a tail of a huge amount of further innovations, building up on the original one. So in some respect even in nature we find the equivalent of these revolutions. In my opinion there still remains the difference that in biology even within these revolutionary times every step itself is as small as possible and 'random' while this is not true for innovations, as there the process of creating novelty involves the human mind and therefore entire (re)combinations of old elements. As we will see this makes some difference when we will try to investigate in more detail how novelty is emerging in the area of economics.

Novelty in the Economic Field

We now try to approach the concept of novelty and it's emergence in the area that we are actually interested in, which is (evolutionary) economics. Whether the various authors prefer a more direct Schumpeterian approach and sticking to the concept of the entrepreneur [8] or if they try to avoid this political motivated concept and treat the topic more in its core [24], it is for certain that innovation - the emergence of new elements - is the driving force behind economic growth. Or at least the only one that has it's origin in the 'real' world, something that leastwise can give us a plausible answer why we should assume that there is growth at all - in contrast to the ever growing level of global debt which seems to assume money coming out of nowhere.

As obvious as it might seem, but first of all we have to answer the question, what novelty actually means in an economic environment. Witt gives a good introduction into this topic:

> Novelty, in the sense of something not previously thought or experienced, is a pervasive feature of science, technology and cultural and economic life. New ideas for acting challenge our ambitions and desires, inventions result in new artifacts, innovations trigger waves of new commercial opportunities and actions. Man-made novelty of this kind fuels competition, structural change, and economic growth.

> These developments, in turn, can generate (often unintended) collective historical outcomes not previously experienced - a kind of novelty of its own. ([24])

Two things become pretty fast very clear among all evolutionary economic authors, and these are:

1. Novelty is by it's very nature unpredictable and connected to uncertainty - A fact that makes it hard to deal with in any theoretical model.
2. Whatever the answer on how to deal with it might be in the end, it is for sure not the neoclassical equation-based approach. Almost every treatment of this topic includes a critique of the standard economics - with good reason.

The first point seems obvious at first glance, but the implications of that statement are much deeper than one might grasp in the beginning. As Witt [25] puts it: "Despite the importance of emerging novelty as an agent of change, a comprehensive, structured theory of novelty does not exist, and it does not even seem clear whether such a theory is feasible at all." This might sound a little bit too negative and dramatic, but nevertheless it is true, especially that there is no structured theory of novelty. There are ideas here and there, but it's unpredictability is what it makes it so hard to deal with. There are some insights from the field of human creativity and cognitive science, but this is of course still no systematic treatment of the topic that could be used for an economic model, or any structured model in general.

As Witt points out correctly, the whole idea of incorporating something unpredictable into our models seems to surpass our mental capabilities. Another problem is that it makes a huge difference if one addresses novelty "ex ante (something unknowable)" or "ex post (not previously known)". While the first approach, theorizing about things before their 'revelation', seems almost impossible from an analytical point of view, the latter on the other hand has the advantage that the observer has already knowledge about what has emerged, but on the disadvantage, that actually there is no novelty anymore in the epistemic sense. Therefore, post-revelation analysis only labels something 'novel' in retrospective. Both approaches lack therefore any real predictive power: "One of the consequences of the epistemic limitations that are binding up to time τ is the predictive weakness of scientific disciplines in whose domains novelty is regularly generated in the course of events, economics being a case in point." ([25])

This *uncertainty* that is connected to the emergence of novelty is also, what makes mainstream economics such an easy target to all its critiques - especially those which are coming from the evolutionary economics side. Witt only uses a side sentence, to already sum it up: "(...) in economics the emergence of novelty is usually treated as an exogenous shock. If novelty is considered at all, the focus is on its diffusion in an industry or the economy." ([25]) Other authors have to say their fair share to this topic too, often from slightly different angles. The following statement from Arthur [2], who was already mentioned in the section about evolutionary economics in general, is also interesting in that regard:

> Behavior creates pattern; and pattern in turn influences behavior. It might be natural in such a setting for economic theorists to study the unfolding of patterns that economic agents create. But this obviously is complicated. And therefore to seek analytical solutions, historically economics chose to simplify its questions. It asked instead what behavior caused an outcome or pattern that leads to no incentive to change that behavior. ([2])

As already mentioned, he also describes the phenomenon of perpetual novelty, where "nothing ever settles down." He describes a game in which simple strategies like tit-for-tat are fighting each other in an more or less infinite tournament. These strategies can evolve by deepening and remembering more rounds of the past. Sometimes, more simple strategies dominated, then more complex, then again simple ones and so on. There was no equilibrium but perpetual novelty. "This is unfamiliar to us in standard economics. Yet there is realism about such dynamics with its unpredictable, emergent and complicated set of strategies." ([2]) Especially interesting is his assessment on when perpetual novelty arises, which could be said - from a provoking point of view - to be true in the real world all the time in every economic aspect:

> There is no precise rule, but broadly speaking perpetual novelty arises in two circumstances. One is where there is frustration (to use a physics term) in the system. Roughly this means that it is not possible to satisfy the needs of all the agents (or elements) at the same time and that these jostle continually to have their needs fulfilled. The other is where exploration is allowed and learning can deepen infinitely - can see better and better into the system it is trying to understand. ([2])

Another approach of critique is to put more the Schumpeterian aspect into the spotlight, which is for example done by Foster and Metcalfe [8]. They are focusing on the fact, that uncertainty results in problems for the *homo economicus*, the rational agent of the neoclassical theory, as the role of the entrepreneur in taking a lot of risk does not seem to be the rational choice in these models, although in reality it could make perfect sense. But one thing after another, let's first see what they say about standard economics regarding uncertainty:

> It cannot approximate economic decision-making when there is uncertainty, i.e., the absence of knowledge of the full set of events faced and the probabilities associated with them. This is the typical state in which technological, organizational and institutional changes occur and these changes, in turn, can create new uncertainties in an economic system. The presence of uncertainty does not prevent economic behavior from occurring. On the contrary, we observe much creative, cooperative and

> competitive behavior in states of uncertainty and the result is 'economic evolution' which is characterized by increases in organized complexity in economic systems and accompanying rises in wealth and per capita income. ([8])

Referring to Harper and Endres, they state that "(...) ignoring the process of emergence in economic systems seriously handicaps economic science" ([8]) and that this emergence has three additional properties, which are: genuine novelty, unpredictability in principle and irreducibility. Genuine novelty is characterized by the formation of radically new bundles of rules, and the enactment of these rules is done by an unpredictable process of *self-organziation*. The unpredictability "(...) is diminished by a process of 'competitive selection' whereby particular combinations of technologies, organizational structures, institutions and procedures come to dominate." ([8])

So here we have again the combination of old elements, combined with a race for dominance among them - classical evolutionary theory. Foster and Metcalfe also mention, in the path of Schumpeter's footsteps, the difference between *strong emergence* which is the main driver for economic evolution and *incremental innovation*, which conforms to the *weak emergence*. They connect the strong emergence mainly with Schumpeter's entrepreneurship and "experimentation with new combinations of rule-bundles". Like others before, they criticize that the treatment of economic emergence has often been incomplete. However, what I personally find most interesting to take away from their treatment of their topic is, as already mentioned before, the question regarding the rationality of the entrepreneur. With the following quotes I try to capture their main message, even though they might be taken a little bit out of context:

> But here we have a paradox. Entrepreneurship is economic behavior par excellence but the entrepreneur cannot behave in the manner of homo economicus, engaging in well-defined constrained optimization. This is because, like an artist or an inventor, to some degree, s/he faces radical uncertainty which means that the totality of possible outcomes is unknown and the probabilities associated with those that are known are unknown. (...) If a decision-maker is trying to operate in a set of interconnected complex systems and is subject to radical uncertainty concerning the potential of the novel project that is being pursued, it would be quite irrational to try to engage in neoclassical optimization, based upon the fragmentary and incomplete 'objective' knowledge available. (...) In other words, it is rational from an evolutionary perspective for a large firm to invest in speculative R&D; to place bets even if most of them turn out to be worthless. ([8])

I already started to tackle the issue of differences between biological and economic evolution and the emergence of novelty within them in the last sub-section. We are now coming back to the question what the 'evolutionary perspective' means regarding economics. As we have seen, the human mind plays a big role in creating innovations and therefore combinations of old elements seem to provide bigger changes (with a certain intention included) than random genetic mutations which are more or less the smallest imaginable changes and only sum up over a long period of time to the creation of new species. However, as well as in economy as in biology there are certain waves of change, which we also could call revolutions. As soon as the temporary stability of the system is not sustainable anymore, one change leads to another - in an exponential way.

Hanappi [13] draws a picture of these dynamics and differentiates between a slow and a fast process (see figure 3.2 for the dynamics in a social context). The slow process is the 'overall process' which generates new regimes as well as the revolutions between them. Within each regime we have the fast process, basically a stable growth process. Only this process can be described by neoclassical theories. But, after some point, this process is not feasible to maintain anymore and revolutions happen. This is the point where standard economics fails to describe what is happening at all. Of course this is also a clever way to incorporate the old neoclassical approaches into the new evolutionary theory by saying: Look, your view was not completely wrong, it just described only a temporary phase of stability. Theses phases exist most of the time, so you were right to some extent. But now the new theory will also be capable of explaining the revolutionary dynamics in between, when one regime ends and a new one starts.

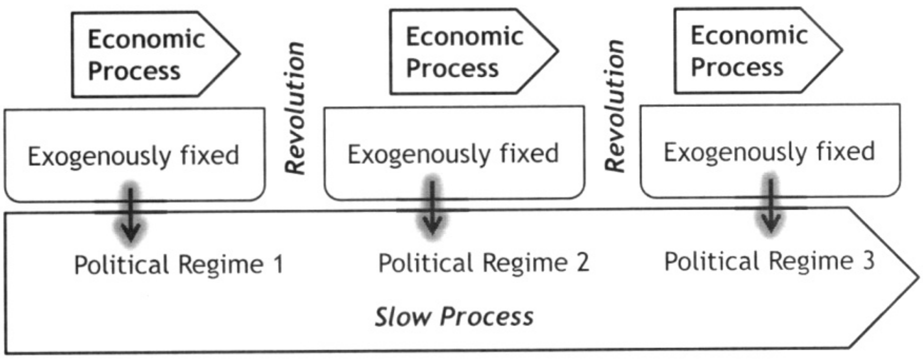

Figure 3.2: Social dynamics of the human species - Source: Hanappi [13]

As already mentioned, these revolutions are in the economic environment strongly connected to phases of mass innovation. Often one technological breakthrough opens the door for dozens of follow-up innovations - refinements of the original innovation or combinations of older elements with the newly discovered ones. But, as in biology, not all novelty is emerging only in revolutionary times. We already said that the fast process is some sort of stable growth process. And growth has to have a source, which is in this case mostly 'weak emergence' - incremental innovations in the economic sense or steady genetic mutation in the biologic sense. Also this 'un-revolutionary' mutation can lead to new species, although it might be much harder to distinguish when the old species ended and the new one began.

How Novelty emerges - New combinations

So we learned now, that novelty is rather hard to treat, and a comprehensive theory could even be completely impossible because of the uncertainty connected with it. That's not the best outlook to be honest. However, all these rather discouraging views on the topic should not let us give up. First of all we have to deal with it anyway, as it is happening in the real world and we already pointed out it's importance. Secondly, even if we don't have a comprehensive theory, we have some ideas how novelty emerges - and these are good enough as a starting point for the model implemented in this work. The keyword is, as the title of this section already gives away, new combinations. But what does that mean?

Well, we already faced Schumpeter's ideas, so the concept of recombining old elements to create something new, should be quite familiar to us. However, aside from rather pseudo-psychological treatments of the heroism of entrepreneurship, Schumpeter, who anyway never created any analytical model, did not take the next step and try to explore how this combination could take place. The same cannot be said about Witt [24, 25]. Regarding the act of creation he refers to Arthur Koestler:

> By a more or less accidental displacement of attention to something not previously noted, he claims, new elements enter into a given context with which they have not previously been associated. Although they come from a different frame of reference, the new elements are in some way similar to the familiar elements, so that a cross-over or recombination becomes feasible. Identifying a previously hidden analogy - an act which Koestler calls "bisociation" - gives rise to the discovery of novelty. ([24])

This already gives us the hint, that just randomly combining old elements will not lead us to any innovation. There has to exist familiarity or at least some kind of connection with the context to which the element is brought, so that this element, which is coming from a different context, still has some meaning. "Only if the conceptual inputs to be blended have some kind of similarity - some common element - can their non-common elements be integrated in an associative act into a new, meaningful concept." ([24]) The similarity criterion though is a matter of conceptual judgment, often ensured by restricting the input sets to be sufficient similar.

One interesting statement by Witt is, that the generative operation "can, in principle, be done by any arbitrary device." Finding sense in these new combinations or aligning the multiple contexts that have been mixed, however, is something completely else and can, from a current point of view, only be done by a human agent. It is also interesting to note, that these two operations in reality will happen in the human mind more or less at the same time. But analytically they can be separated, which is especially interesting if you really would like to implement a generative engine. To put it more formally, here is Witt's first hypothesis:

Hypothesis 1 *The creation of new cognitive concepts (ideas, imaginings) involves two operations. One is a generative operation that produces new (re-) combinations of elements. The other is an interpretative operation by which the new (re-) combination is integrated into a newly emerging or a more general already existing concept.*

([25])

CHAPTER 4

How to model Novelty

4.1 Witt - Propositional Networks

Now the border between the theoretical investigation and a first attempt to put these theories in some sort of model or formal treatment becomes blurred. Therefore we will continue here straight on with Witt's thoughts from the last chapter, even though there will be further theoretical considerations too. He starts with describing economic activities as simple factual statements. These statements only consist of subject S, predicate P and object O. “Batman fights villains.” would be such an example. But let's look at the example given by Witt, which are the two statements “hair (S_1) grows (P_a) in tufts (O_a)” and “hair (S_1) can be cut (P_b) by clippers (O_{1b})”. As we can see, these two statements share the same subject (hair). Borrowing from cognitive psychology, Witt introduces the concept of *propositional networks*, which is used to describe the sharing of a common concept. Formally this can be pictured like in figure 4.1.

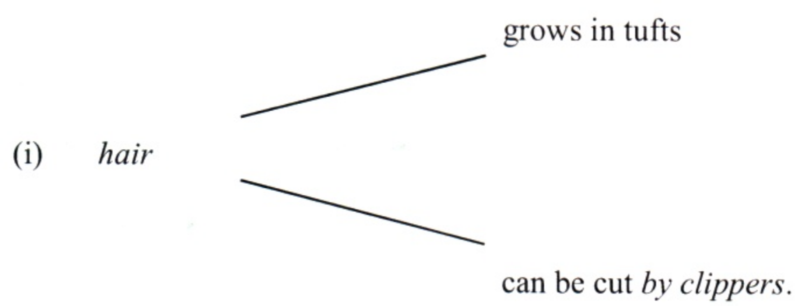

Figure 4.1: Propositional network (i) - Source: Witt [24]

This is obviously a simple and incomplete example, but in general if people stumble across a concept like “hair”, in their mind they automatically connect it with other things, coming from their memory and so on. If they are given another concept, for example "grain", then maybe the following two associations are the first that come to their mind: "grain (S_2) grows (P_a) in tufts (O_a)" and "grain (S_2) can be cut (P_b) by scythes (O_{2b})". See figure 4.2 for the corresponding propositional network. It is rather obvious, that these two networks share the same syntactic structure and some common concepts, which we consequently named the same (P_a, O_a and P_b).

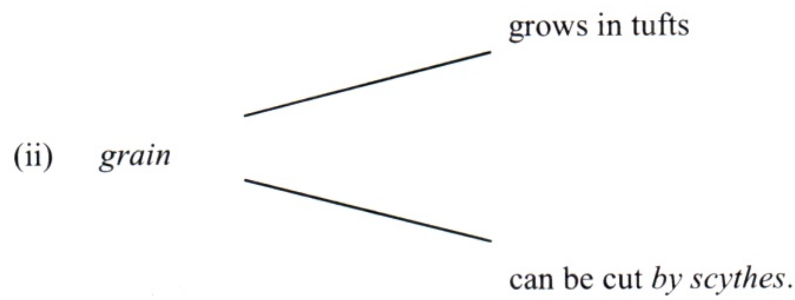

Figure 4.2: Propositional network (ii) - Source: Witt [24]

Of course this extreme similarity and simplicity is forced for the purpose of a showcase and would not be found in reality, but it is not unimaginable that shared elements between concepts like these would occur in a more complex, but similar way. As these two networks are so similar, it is tempting to exchange the subjects and see, if the new statements would still make sense. As the first statement is the same anyway for both subjects, our two new statements would be "hair can be cut by scythes" and "grain can be cut by clippers". Now we see our problem of modeling this whole process, because the interpretation of these new combinations is totally dependent on the intuition of a human individual. It would be commonly acceptable to say that the first statement does not make sense at all. With hasty judgment one could say the same about the second, as cutting grain with scissors seems, while being possible for sure, rather inefficient. At least as long as you are thinking in the way of scissors of the size which are used for hair. A more creative mind could start to think other scales - which "is said to have inspired McCormack's invention of the mechanical reaper (Martindale 1999)." ([24])

Obviously, this does not make the original statement invalid - grain can still be cut by scythes. Therefore a new statement is created, which is called a *combinatory extension* and has the following form: "grain (S_2) can be cut (P_b) by scythes and clippers (O_{1b+2b})". So, in generally we see that the *generative operation* which puts combinations together, does not need semantic knowledge and could therefore be carried out by a machine. Nevertheless, to avoid nonsensical statements, it is not recommended to reshuffle elements just randomly. What is effective on the other hand is to use the commonalities, the similar elements, as a starting point for crossing over the non-common elements.

Hypothesis 2 *The generative operation is based on a contextual pre-selection of sets and elements of sets to be recombined that follows some similarity or confor-*

mity criterion.

([25])

Therefore, the generative operation is far from a "blind" variation, but there is a strong guidance towards what is chose to be recombined. This means, that if we want to use analogies to biology, the term of *sexual selection* would be more appropriate than blind genetic mutation. Anyway, the generative operation is actually the far less troubling part of the problem. The real difficulties lie in the *interpretative operation*. It is the operation which identifies sense and semantic meaning in the new combinations. A machine obviously cannot do that. It seems like a human observer is at least necessary for this part, after the combinatory extension are artificially created. "Unlike the generative operation with its plain logic, the associative act by which the human mind accomplishes a conceptual integration in the interpretative operation is not fully understood." ([25]) As Witt puts it, we by-pass this problem by "letting our intuition doing the trick." Hypothesis 3 just sums up, what I already indicated:

> **Hypothesis 3** *While the generative operation can be automatized mechanically or electronically outside the human mind, for example, by numerical algorithms and programs, the carrying out of the interpretative operation is bound to the medium of the human mind and can therefore not be automatized.*
>
> ([25])

At this point, a clarification is probably needed: The interpretative operation is not identical with the evaluative one. In fact, these are two different operations. The first one only answers **what** it is that emerges, but it doesn't say anything about the utility or benefit. This is obviously done by the evaluation. In general, the evaluation is not necessary for creating novelty per se. However, it can be included in a feedback loop which then determines the criteria for the preselection of the generative operation. Interestingly enough, the evaluation step can be automated again, as it is logically distinct from the interpretation. But even if we know the algorithm of the generative operation and can automatize the evaluative one, it is usually not possible to predict the implications from the premises. This is because the recursive connection between our three operations are too complex, and therefore experiments and simulations are necessary:

> **Hypothesis 4** *Novelty is revealed through inductive operations whose outcome, by definition, cannot be derived from the premises. Carrying out the inductive operations requires time, and thus prevents the meaning of novelty being instantaneously accessible.*
>
> ([25])

In reality, often even the generative operation can be chaotic or too complex to be solved analytically. But even if there would be a possibility to automatize the interpretative operation in a non-trivial way, the time constraint would still hold because of the complex recursive relations between these operations. It might be obvious, but still has to be stated: There is and will be no way to go through all steps until a certain point of technological level only in a theoretical way and therefore skipping all steps in between - we have to inductively go through all steps by carrying them out in reality. This also means, that the complexity of these relations will be probably not getting lower, but on the opposite rise with every new element as we most likely have an auto-catalytic process here:

> **Hypothesis 5** *By re-using newly created concepts in further iterations of the generative operation, an infinitely growing number of concepts can emerge from a finite number of initial elements, provided that the share of combinations to which new meaning can be attributed is non-vanishing.*
>
> ([25])

Hypothesis 5 brings up some very interesting thoughts. First of all the condition mentioned at the end: What would be, if there is some kind of end-state? Where new combinations do not result in meaningful novelty anymore? Right now is hard to imagine, on the other hand there is some kind of starvation regarding big innovations to be experienced currently, although it probably has more economic than technological reasons. Also, one would think that with increasing complexity and longer chains of concepts composed, there could be some point where the human mind would not be able to handle that anymore. This is not the case because there is a complexity-reducing habit of denoting longer chains by introducing new shorthand concepts, like ‘supercenter’ for the hybrid of a kiosk and a gas station. This new concept can then be used for new combinations without substituting any of the already known elements which are only parts of it.

Another interesting thought is, that if you use backward induction and following the logic of this hypothesis, you have to arrive to some point where there exist basic elements which can not be derived from others anymore. They are called *innate concepts* according to the *embodied knowledge* hypothesis. In general, the assumption of an *autocatalytic* expansion seems to make sense when we look at our own history and the accumulation of human knowledge. But, as already mentioned, it is possible that there is a limit. Maybe at some point there is no new knowledge to gain anymore. Maybe, at some point we are just not able anymore to make sense of new combinations or make them available for further (re)-combinations because of memory and communication constraints.

Further, Witt raises the intriguing question if there exists some sort of degree of novelty. As already mentioned, there seems to be the difference in economic evolution compared to the biologic one, that novelty does not have to occur in the smallest possible steps, but that there are some leap jumps. It does seem to make a difference if the new product only changes a minor detail of something already known (making it nevertheless novel, but in a rather 'unspectacular' way) or if it is something so entirely new that many people wouldn't even have imagined it before to be possible at all. It is possible to distinguish between basic and follow-up innovations - at least **ex post**. But how about **ex ante**? What we can say rather sure is that a higher degree of novelty is connected to more novelty-induced uncertainty. Also certain constraints on the generative operation will influence the possible degree of novelty. If the rules how to generate new elements are rather strict, big surprises are very unlikely to happen. Regarding the interpretative operation, it is rather hard to find such constraints, as the operation itself is not really understood. It seems that it is either non-trivial and therefore without any constraints or trivial and the meaning of all possible outcomes is already known before, so no associative act of the human mind is needed anymore:

> **Hypothesis 6** *The ex ante degree of novelty is minimal (but not necessarily zero) if the generative operation is known and the interpretative operation is trivial in the sense that the meaning of all its possible outcomes is known beforehand. The degree of novelty is maximal where the generative operation is not known and no constraint can be imposed on the outcome of the interpretative operation.*
>
> ([25])

While this might be a true and a good basis for a more formal treatment of the topic, it is at the same time weak and vague. What does it mean for all the degrees of novelty between those two extremes? Also, what does it say about uncertainty in the extreme case of complete knowledge? Let's start with the latter question: One example for minimal novelty would be the result of a lottery: As we know the drawing mechanism we also know every possible outcome. But, it is a random mechanism and therefore the outcome is still uncertain.

However as we know all possible outcomes we can calculate probabilities. The question is if in the space between minimal and maximal uncertainty this can also be displayed in a probabilistic way. Unfortunately, the answer is no. In both cases - if the outcome of the interpretation is non-trivial but the possible outcomes of the generation are known, or if the interpretation is trivial but the outcomes of the generation are unknown - there is no possibility to represent this uncertainty in a probabilistic way. Therefore all we can say ex ante is if the degree of novelty is minimal, maximal or something in between. "Because of its peculiar epistemological status, novelty-induced uncertainty is a much less structured concept than uncertainty in the probabilistic sense." ([25])

4.2 Hanappi - Evolutionary Selection and Variation

Another sort of basis for a modeling attempt that we are looking at comes from Hanappi and Hanappi-Egger [14]. It's coming from a completely different direction with a more biological view on evolution (while being still an economic paper). While there are not much intersections with Witt's theories (besides new combinations of old elements, obviously) it is important to see also some approach regarding the generative operation and the selection mechanism - the latter is something which Witt only shortly mentioned as evaluative operation, if he even means exactly the same at all.

Hanappi and Hanappi-Egger relate in their paper very strongly to Schumpeter, but more in the sense of (re-)combination than in any political regard, and obviously their ideas are very similar to those in this work here and one of the sources of inspiration, which should be rather less surprising as Hanappi is the advisor of this work:

> The idea that in the course of time elements from different old overall processes might be copied and combined to form new overall processes adds dynamics to the usually static set theoretic framework. Such a theory of development aims at two goals: It should be able to explain how old combinations came about in the past in the first place, and it should suggest some hypothesis how and where contemporary new combinations will emerge. ([14])

Obviously, as we already explored, the second goal seems very hard to achieve - especially the 'where' could be even impossible in concrete cases. But let's have a more general look at the selection process first. Following Darwin, we have "(...) a set of individual members with different characteristics and an environment that in the course of time deletes members with unfavorable characteristics." ([14]) If we have n different types of species entering the environment and the filtering characteristic of this environment is so strong (strict binary *survival gate*) that only one species survives, while the carrying capacity is equal to the number of entering species, than the invasion process (selection and breeding) looks like in figure 4.3.

The growth process happens at a speed of doubling per period and as the environment is stable, the survival gate remains at the same position all the time, which means that all selection already happens in the first period. Only when the carrying capacity is reached, individuals of this perfectly fitting species are also starting to die - they are competing for the same resources. From then on there is obviously no growth rate anymore - and as we have a breeding speed which is doubling the population at each period, half of the species has to die.

Obviously this is a pretty simple scenario, and also the authors are referencing to Simon [22]: "The simplest scheme of evolution is one that depends on two processes: a generator and a test. The task of the generator is to produce variety, new forms that have not existed previously, whereas the task of the test is to cull out the newly generated forms so that only those that are well fitted to the environment will survive." In this simple setup we only have a real generator of new forms in the beginning by the introduction of the different species, afterwards it's only a replication of existing ones. However even this simple setup leads to some interesting thoughts, of which I want to quote the most interesting one:

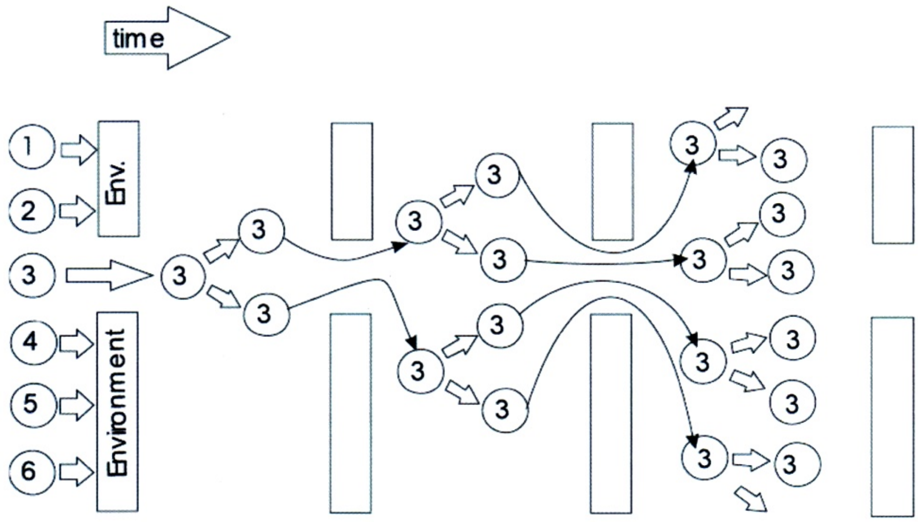

Figure 4.3: Selection in a stable environment - Source: Hanappi and Hanappi-Egger [14]

> Consider the idea that the variety of types invading the environment might be variable. If the window of survival (type in the example) is unknown to invaders, then the probability of invaders consisting only of one type to survive period 1 is very low. To be precise, with a window of size 1 and n possible type it would be $\frac{1}{n}$, whereas m invading types, $m \leq n$, would have a chance of $\frac{m}{n}$ that one of the passes to period 2. If survival in period 1 counts as a necessary condition for success of a species, then variety of types certainly contributes to it: There is a positive relationship between the variety of invaders and the survival probability of the species they belong to. ([14])

One might wonder how this is related to the model of this work and also the topic of emergence of novelty via new combinations in general. Without spoiling too much, if we are imagining the 'invasion' as the introduction of new products and the survival gate as some kind of market feedback, this is indeed an interesting starting point for modeling novelty also in an economic environment. But further considerations are necessary for a serious application in the area of novelty simulation. First of all there is the question, where these species with variety in the beginning came from. One plausible answer is that the environment is only stable for a certain time - a long time compared to our reproduction process, but nevertheless there are revolutionary times in between. This is a concept that sounds quite familiar, which is not surprising as we introduced this view of evolutionary dynamics by Hanappi already in the previous section.

Without the possibility of going too much into the details of the work of 14, as it would go beyond the scope and constraints of this master thesis, I will try to summarize the most important thoughts regarding possible extensions of this model. One thing that comes obviously immediately into the mind is letting the survival gate change over time. In our case that would mean the immediate death of the whole uniform population. As this seems rather unrealistic, the real question is therefore, how variation within a species can emerge? For the authors, genetic mutation is not an satisfying answer as the chances of a favorable mutation in one single gene would be one to a billion. Therefore this type of improvement would be way too slow to explain what happened on earth.

Another source of variation would be regarding the *adaption speed*. If the environment changes in some periodicity (day and night, seasons, ice ages and so on) and certain species have the characteristic to change in similar oscillations, they would be favored. Due to the selection forces of evolution, the periodicity in the environment produces a structure of innate frequencies in the members of the species. Personally, I doubt that there will be found too much of these oscillations in reality despite the already mentioned - neither in biology nor in economics. In my opinion, the sudden changes, the shocks and catastrophes are the real driving forces of novelty.

However, in both cases there is the question regarding the possibility of anticipation, which according to the authors also exists in evolution. In the case of periodicity an anticipating species would then switch from one variation type to another, namely the one it anticipates has the best fit to the new survival window. The authors therefore assume some kind of additional consciousness which is necessary for making a complete switch from one type to another possible and also the properties of the currently unfavorable types have to be stored somewhere for later. While there could be doubts about the implementation of this principle in the biologic field, it works perfectly fine for economics.

Without going too much into the details of the authors model, for their findings they experimented with a species consisting of six different types, a survival gate which lets pass two different types in one phase and they varied the speed of change of this gate, as well as the innate growth rates. One of the main conclusions is that only if the environment changes every period, a maximal *diversity drive* of 6 is optimal. In slower changing environments a low diversity drive leads to much higher survival rates. This is not very surprising, given the setting of the model. Further there is a negative relationship between instantaneous diversity and speed of environmental change. That suggests that faster changing environments will select more diverse populations.

Regarding the innate growth rate there is one interesting result: Only in slowly changing environments the optimal drive to diversity increases with the growth rate. That means, the higher the growth rate, the better it is for the species in a slowly changing environment to have a higher diversity push, to redistribute the composition of the different types within the species. One explanation taken from reality could be that otherwise quickly too many exemplars of the same type would exist, which could not as easy coexist as individuals of different types (less fight for common resources, different enemies of other species, and so on).

Obviously the model is missing something important for our use: The emergence of completely new types. The authors come up with the idea of **genetic algorithms** as a solution: “In a genetic algorithm the behavior of a type is represented as a bit string, and one could think of the innate growth rate of this type as its fitness function. The crucial assumption then is that parts of the bit strings of successful types can be glued together - and the new fitness of this newly born type is immediately available.” ([14]) Now we are already coming very close to what the model of this work is supposed to do, although I will use vectors of product properties instead of bit strings to describe the fitness of the new production processes which will be derived via combinations of old ones. For further details I have to refer to the next chapter though.

As the authors are referring to Schumpeter as a starting point in their paper, they come back to him now at the end: According to them, random mutation would be pointless, there needs to be an entity which already preselects the elements for combination before the fitness of the result is known - this entity is the entrepreneur. They also introduce the term ‘vision’, something which in economic and technological innovations is obviously necessary in reality, as it happens rather seldom nowadays that somebody stumbles accidentally over novelty. Therefore they argue that if we want to find analogies in the emergence of novelty between economics and biology, we have to look at sexual selection, where there is not only randomness but some sort of intention. While I would agree with the last part of this statement, I am still not sure if sexual selection is the right process to compare with; after all the selection criteria there are often rather unconscious and irrational and not with the intention to create novelty but with the purpose on replicating the positive attributes of the selected partner.

4.3 Kauffman - Peeking at Biology

Even though there was already a lot of usage of biological terms and ideas, the previous model was still described and built with the intention to use it in an economic environment. Now we are going completely into the field of biology and will see what they, or respectively Kauffman [16, 17] as an important contributor, have to say regarding the emergence of novelty. Well, first of all we have to realize that we are in a completely different domain now. Kauffman is not talking about recombinations per se and also his conclusions which are quite revolutionary for biologic evolution (and for which therefore obviously also exists some criticism and reserved judgment from colleagues in the field) not necessarily have to translate to economics. Nevertheless his ideas and models are interesting per se and can be an inspiring source for modeling in the area of novelty, which is why a look at it can broaden the horizon for sure.

The main premise of his theory is *self-organization*. According to him the emergence of life was not a lucky and unlikely event but rather an unavoidable one. Also it was not a slow progress, based on pure random mutation, starting in a very simple state. Kauffman believes that life was never simple but always complex. The reason for this is that according to him in sufficiently complex non-equilibrium chemical systems like the one we are living in, there always emerge *autocatalytic reactions*. What does that mean? In short, summarized by a biological layman, it means that chemical reactions happen much faster if there is a specific catalyst present. The system becomes autocatalytic when it produces in some of it's reactions elements, which are working themselves then again as catalysts in these reactions. It is pretty plausible that in such a setting once the autocatalytic process starts, a chain reaction of catalyzed reactions will crystallize.

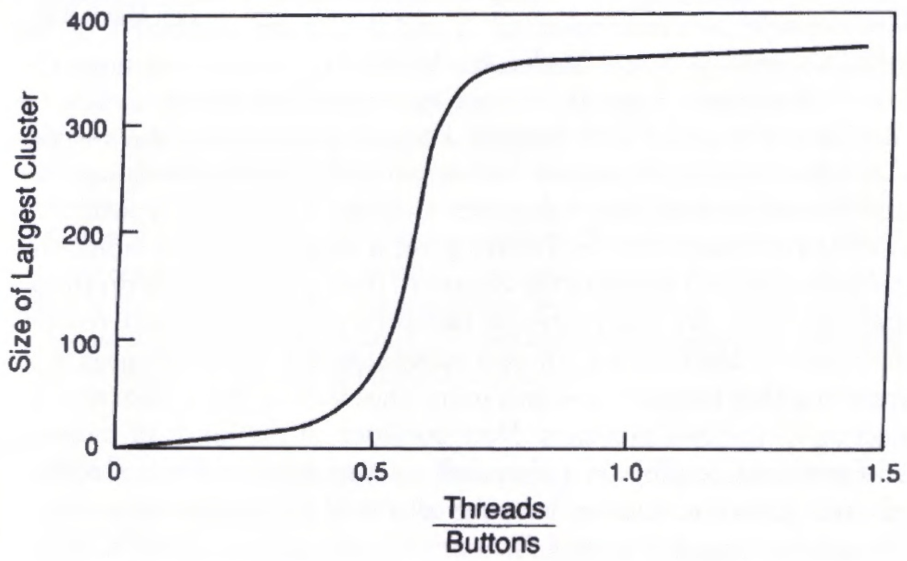

Figure 4.4: As the ratio of lines to edges passes 0.5 the size of connected clusters is exploding during the phase transition and a giant component crystallizes. - Source: Kauffman [17]

A simplifying model of this process is the one of a *random graph*. If we have a predefined amount of nodes (let it be buttons for the purpose of a more figurative example) 'lying around' and we start to randomly connect this nodes with lines (threads) then at some threshold of threads, which is around half of the buttons, the buttons get so strongly connected to each other that if we would pick up one of them, a whole web of buttons and threads would follow - only few buttons would still be unconnected. Kauffman is talking about a *phase transition* (see figure 4.4). Simplified this means that the change (and in our case the emergence of life) did not happen in a slow steady development but rather in one big jump. The bigger the system, the more vertical is the steep part of the *sigmoidal curve*. If there were infinite nodes than there would be even a real discontinuous jump - a binary transition from off to on. In the biologic field according to Kauffman such a web is almost certainly autocatalytic and therefore self-sustaining and alive.

The main thing to stress here is the difference between *spontaneous reactions*, which happen very slowly and these catalyzed reactions, which are assumed to occur rapidly. When we look at figure 4.5 we see, that this reaction network can be easily interpreted in our sense of combinations of old elements, although in our case the results would be probably not that predictable. However, there is one big difference: Until now we only differentiated between random, acci-

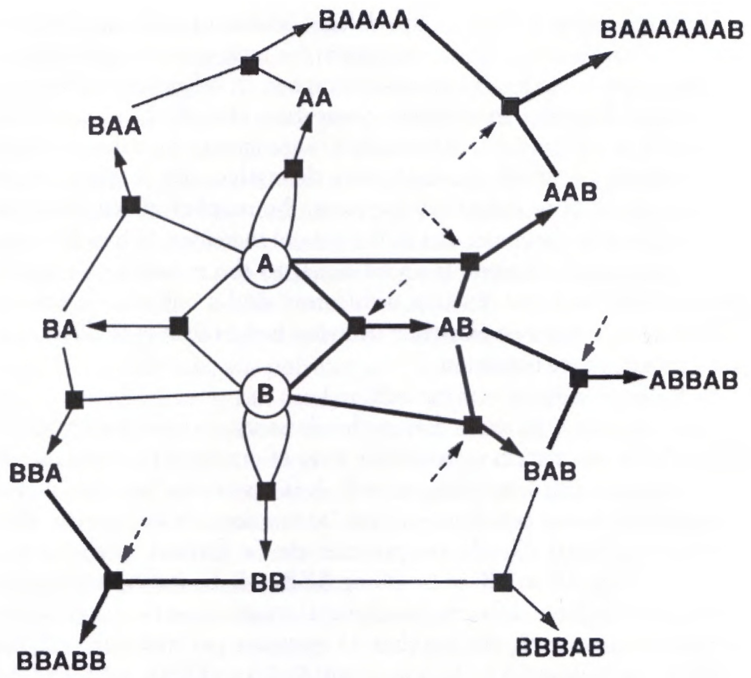

Figure 4.5: Network of reactions - The reaction squares indicated by dashed-line arrows are catalyzed, and the heavy darker lines connect substrates and products whose reactions are catalyzed. The result is a pattern of heavy lines indicating a catalyzed subgraph of the reaction graph. - Source Kauffman [17]

dental mutation (of which we already declared, that this is usually not the way novelty emerges in the economic sense) and intentional recombination, which is most comparable to sexual selection in the biologic sphere.

Nevertheless, this is still in the sense of spontaneous reactions, as it is a very slow process. The idea of catalysts or even an autocatalytic system is new to us in the course of this work. According to Kauffman, the emergence of autocatalytic sets is almost inevitable. The question for us is, how far are they relevant for our emergence of novelty? We know from our own history that there are phases, were innovation happens everything but slow. Key innovations lead to chain reactions and tons of extensional innovations. However, personally I find it hard in the economic sense to determine if such a key innovation should be considered as 'input factor' or as catalyst, and probably this has to be defined from case to case.

Another point that Kauffman covers is the question regarding evolution itself. Until now the main question that was tackled was how life could emerge at all. But then another question arises: How can there be selection and variation before the DNA and genetic code was 'introduced'? Especially as according to our picture of evolution, such an evolving system has to be basically stable, but still be able to change in slight alterations. The autocatalytic systems we discussed so far though are most likely pretty chaotic and unstable.

For his explanation, Kauffman uses another mental picture, another idealization. Enzymes can be 'turned on and off' by other molecules. This is not true as there is graded catalytic activity and not a boolean one. However, the change is often very sharp and can therefore be seen as having two states. Also, for this picture the substrates and products are either present or absent. This is of course also not literally true but the concentrations often change swiftly from low to high. The picture of the metabolic network of enzymes, substrates and products is now an electrical circuit of connected lightbulbs. The bulbs influence each other in their behavior of blinking. The question is now: Can their arise an orderly, systematic behavior of blinking on its own or does it have to be complete chaos without careful crafting of this system?

As we can address each state of a bulb with on or off, and as also the bulbs are influencing each other, we can easily use *boolean algebra*. For example two bulbs can determine via the 'AND' or the 'OR' rule the state of a third bulb in the next time period. In a scenario of three bulbs, where every bulb influence the others, we can easily show for every state in T how the states in $T + 1$ will look like and find the *state cycles* in which the system will fall into pretty quickly, depending on the initial state. An noteworthy property of *boolean networks* is that while every of them settles down to a state cycle, these cycles can vary drastically in size; i.e. the number of states it cycles through. It can be stuck in one single steady state or cycle through such a huge amount of states that there is no order recognizable anymore. As a small side remark, the concept of the state space described here, even though coming from a completely different direction, is not completely different from what we already briefly discussed in the model of Potts (see chapter 3).

So a system that qualifies for following the evolutionary principles should have rather short, but not too short state cycles. Kauffman introduces the concept of *attractors*; state cycles where the system flows into it after some time. It does so because these attractors have some kind of gravitational basin surrounding them, that leads into them. We can imagine the attractor as a lake and the basin of attraction as the water drainage flowing into it. A lot of different trajectories can settle into the same attractor, but several attractors exist at the same time, and often the initial state is deciding in which attractor it will arrive. So for a system which is stable, but not completely rigid, we need tiny attractors that 'trap' the system in tiny subregions.

Another prerequisite is *homeostasis*. That means, the system must be resistant to small disturbances. If for example a single bulb is randomly chosen and flipped to the other state, it should leave the system in the same basin of attraction. Otherwise even small changes would turn the system upside down. The same is true if we rewire the network and for example change the rule of a bulb from OR to AND. Kauffman suggests, that stable and orderly systems in that regard also evolved automatically - he calls it "order for free".

Two features are necessary to determine if a system is ordered, chaotic or on a phase transition between these, "on the edge of chaos". One is the number of 'inputs' for each bulb. The other is the bias regarding the control rules themselves, like OR and AND. Some of these rules tend to create order, others are doing the opposite. Regarding the number of inputs - one could also call it connectivity (K) - it is easy to see that the lower it is, the more order we get. With $K = 1$ there will be very quickly a short and ‘boring’ state cycle, often consisting of only a single state. On the other extreme, with $K = N$, which means every bulb gets inputs from all other bulbs, we have 2^N possible states - pure chaos, even though the number of attractors is much smaller.

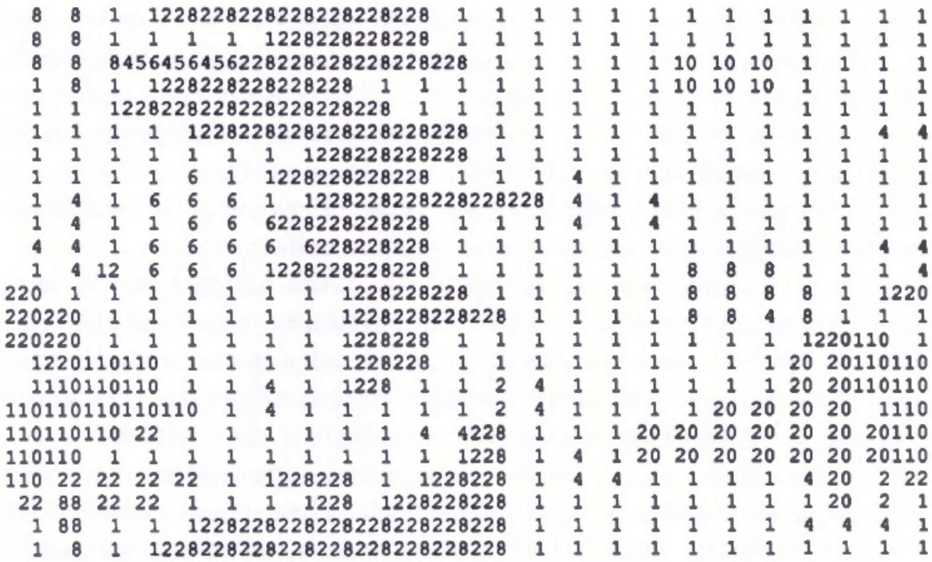

Figure 4.6: Two-dimensional lattice of bulbs. The numbers indicate the cycling period. Therefore all bulbs with 1 are red and frozen, the others are green and switching between on and off according to their cycling period. Source: Kauffman [17]

According to Kauffman $K = 2$ is the magic number. For these networks, the length of the state cycles is around the square root of the number of binary variables. For example a network with $N = 100000$ bulbs would have 2^{100000} possible states, but cycles among only 317 of them. For evolutionary principles we seek exactly these systems which are in order, but on the edge of chaos. “Just between, just near this phase transition, just at the edge of chaos, the most complex behaviors can occur - orderly enough to ensure stability, yet full of flexibility and surprise. Indeed, this is what we mean by complexity.” ([17])

Another simplified picture is necessary: Let’s imagine a two-dimensional field of bulbs (see figure 4.6). There are bulbs which are frozen at on (1) or off (0) from the beginning, as they are never getting influenced by the others - these bulbs are red. The other bulbs are green and can switch between on and off. In a completely orderly system everything would be red, and we would have no dynamics. In complete chaos everything is green and just blinking seemingly random, everything reacting to everything according the *butterfly effect*. In both cases, there is no possibility to reliably send a signal from one side of the lattice to the other. Only on the edge of chaos we could have small dynamic islands and some tendrils of contact between them, to make this possible and have some sort of coordinated change without effecting the rest of the system. Kauffman suggests, that evolution could favor and select systems like this.

Naturally the question arises: What does this mean for our economic emergence of novelty? Does it mean anything at all? Well, part of the exercise was just to show which possibilities exist in general to model the emergence of novelty and evolution. For sure, the tendency towards self-organization in an autocatalytic way is an interesting aspect, but hard to incorporate in our approach of recombinations, which is coming from a completely different side. As already mentioned, for an incorporation of the (auto)catalytic principle, the first task would be to find a consistent distinction between catalysts and input elements. Also, it is hard to define what the economic analogy to this biologic order would be. The picture of the state space seems to make more sense if we look at economics in a general, theoretical view. One idea is, that this tendency towards order is exactly what we see in our model of social dynamics (see figure 3.2) when stability arises after a chaotic revolution. But now in my model I want to specifically show the emergence of novelty resulting from combinations of old elements, carried out by concrete agents, and agent-based systems are known to have their own ‘ordered’ dynamics. Nevertheless these concepts could be an interesting extension of the model presented in the following chapter, which is obviously the centerpiece of this master thesis. And maybe, this tendency will even be observable to emerge automatically in some way in my simulation runs.

CHAPTER 5

Modeling and Simulation

5.1 A Model of Innovation with re-combinable Production Processes

After our journey through all these relevant existing ideas and literature I hope I could give you a nuanced overview of the topic with views from a lot of different perspectives. Now it is time to introduce the 'core piece' of this master thesis, my own model. To give a quick overview what this model is about again, I will repeat the introduction: The aim of the model is mainly to build a showcase on how a concrete model can use those previously explored theoretical ideas. The model built in this work will focus on the production side, while using the market only as a sort of testing mechanism, to find out how good the new products fit the customers' demands. The innovating agents (or entrepreneurs) have access to a pool (which has not to be the same for every agent) of (partial) production processes. For those processes certain parameters are known from historical experience, like the input in terms of costs and the output in terms of customer utility. This utility is characterized by a vector of quantitative measures for different product properties which can be matched against a similar vector which represents the customer's expectations.

The innovators will try to recombine those production processes to create a new process which represents a new product. A key concept for this model are therefore the links between those processes, which define if a combination is possible at all, and if so, what the already mentioned parameters of the new process look like. When leaving the view of a single agent and zooming out to an aggregated view, this will lead to a dynamic race for innovation, where older products will experience an utility discount over time and the different producers will try to come up with new products, do constant market research and eventually produce successfully tested new products, while being in permanent competition with each other. The questions to be answered are: What information does the simulation provide regarding the emergence of novelty in such an environment? How well do the simulation results match with what we can experience in reality?

How can this model be enhanced or used in other more sophisticated evolutionary models, that need such an innovation engine, but focus on other key topics?

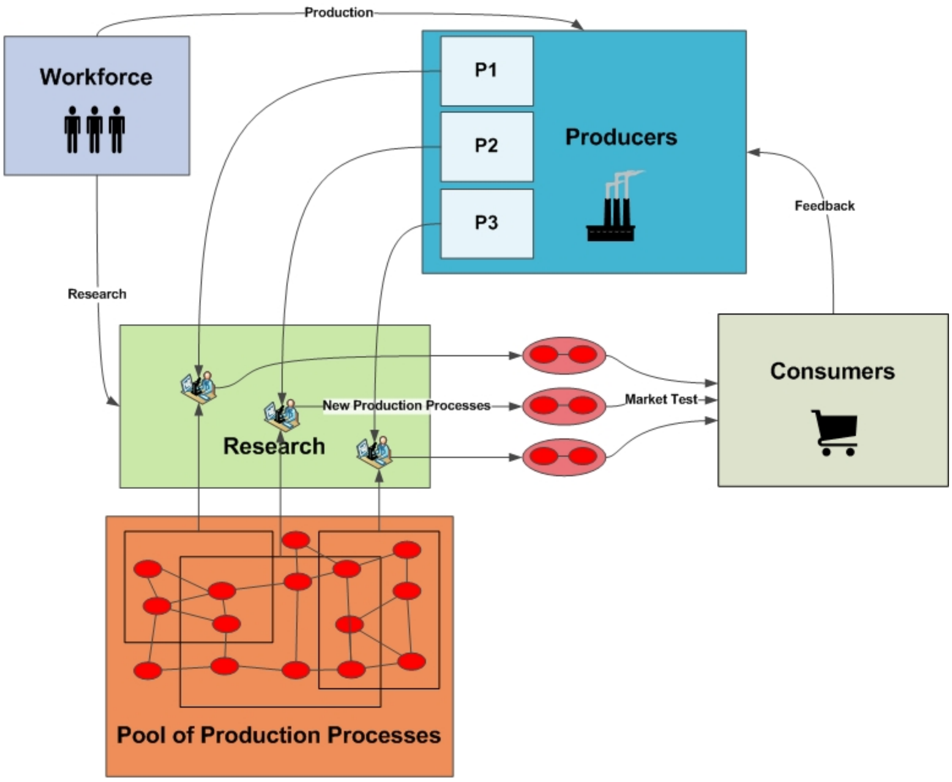

Figure 5.1: Basic structure of the model

Research & Production Process Pool

A first overview of the basic structure of the model you can see in figure 5.1. As the conceptual graphic shows, we have several *producers in competition* to each other. Each of them is performing its own *research*. The inputs for this research are first of all *workforce* and second, and for the focus of our model even more important, the *pool of production processes*. As we can see, these processes are linked to each other. In the figure only a couple of connections are represented, while in our model every process is in some way connected with every else. These connections are equally important as the nodes themselves, as they define the properties of a new process which would arise when one of the researches tries to combine them to a new process.

However, this does not mean, that every combination does make sense. A lot of these potential new combinations have an indicator in their properties that shows that a combination is not possible and therefore not creating any useful new process. So in a way you could also say that there is no real connection then, and while there is in theory still one which just says, that this connection does not produce a new combination, the experience with NetLogo has shown that modeling links and especially whole networks like that is very resource-intensive. So therefore for the purpose of reducing the processing time, the implementation of the model will also spare a technical realization of links which are not useful. Then the non-existence of a technical link exactly means, that the combination of these two old elements does not produce a new one. So even the non-existence of a technical link contains in fact some information.

This whole 'predefined' network might be considered as a drawback of the model, because if you want to see it that way, one could say that there is some static knowledge of all possible combinations even before any researcher tried it out. All the agents then additionally do is to explore this predefined space. This is maybe not the picture we would like to have of innovation: That all we do is searching for solutions which are already there. Probably we would like to think that it is the inventor, who is creating the solution. Therefore I wholeheartedly agree, that here is room for improvement. However, every idea needs a start, and this is good enough for this first attempt. Also one should consider, that the network is not static at all. With the combination of old elements new nodes will be created. Every new node generates new connections with all other nodes - even with those from which it originates.

We can also see, that every company has a different view on the pool of production processes. This makes sense, as companies not only differ in their specialization regarding the products they manufacture, but even within the same product group there might be different strengths and weaknesses among them. It is even imaginable that almost the same product is done with different production processes by competing companies. The question is now: How can these different views change over time? First of all, obviously with the discovery of new combinations, the network of possibilities also for each company changes. But it would be unrealistic to assume that every company is making all discoveries on their own. Therefore they also can learn form each other. After a certain amount of time that a new product (and the underlying production process) is on the market, it becomes common knowledge for all competitors.

Even before this time span is over, it is easier for companies to find out about an combination that others have already discovered before them. In general the producers have the free choice which new combination of old production processes (in our model for the sake of simplicity this will always be limited to two processes) they want to investigate. When they start a research there is a certain probability in each time step (in NetLogo we are working with *discrete time*) for a *breakthrough* - which can also just mean the realization that a combination is not feasible at all.

This probability is dependent on the workforce of researchers, but also on the combination itself. Certain production processes *fit* better together than others. Of course the producers are not without any knowledge either, so their task is to select exactly combinations which are promising success in their view - This means they share some technological connection which suggests that a combination of them is either an obvious choice for improvement or would be exceptionally productive or a big leap forward, even if it could be hard to achieve. Exactly this 'knowledge' or 'instinct' is what a surviving company can distinguish from one which will die in the race, and of course it is also what can put a small startup company ahead of a big concern, even with much less budget and therefore manpower in research. The implementation of different combination strategies is for sure one of the first expansion which could be made to the model, as for now the selection of combinations will be done based on successful experience with certain processes and a lot of randomness. But thanks to the agent-based approach it is an easy task to implement different kinds of agents with whatever strategy should be explored.

As soon as a competitor discovers the combination that another company is researching too at the same time, the latter can continue with a higher probability and therefore usually a faster discovery (a possible extension would also be to implement a decision possibility to stop researching this combination totally, as it wouldn't be technology leader anymore). With every company who has gained the knowledge of a combination, the probability gets higher for all others which are still researching it. For combinations that don't make sense, the assumption for now is that companies would not share the knowledge of failure. This implicates that even after the discovery of this fact, nothing is changing for all other companies which are researching it too. Of course this is also a topic which is for sure a legitimate subject for debate or extensions.

One last remark regarding figure 5.1: In the picture we see production processes which are not in the view of the pool of any company and therefore unknown to society. Now we have two different options of interpretation. The first one is to say that every existing process has to be in the view of at least *one* producer. We only pictured three producers in our schematic depiction, but of course there could be more in the model. So in the implementation (or in reality) some other producer would have it in its view. This interpretation will be our standard variant of the model. The argumentation is, that otherwise it could get hard to explain why there are some production processes which emergence in a different way than with recombination. We all the time said that recombination is the basic principle for us, and then we allow that by random chance new production processes also could be discovered completely out of nowhere?

If we think about the technological development of the last centuries that seems rather unlikely. But in the full course of humanity, maybe there are some foundational discoveries which can't really be described by the incremental approach, at least not in the sense of recombinations. Therefore an extension which allows random discoveries as 'side effects' of research, of course on an extremely low probability, could make sense - even if it just would be to examine if the introduction leads to any differences in the results of our simulation. Again, this is a good suggestion for further extensions of the model.

New Combinations & Market Test

We already talked a lot about how researchers try to recombine old production processes to create new ones. But what attributes belong to these processes? How can you assess if a process is useful or not? And how to find out if the consumers like it or not? Well, first of all the basic input of each process is costs per unit. As we will see, in this model workforce is not playing an important role and therefore abstracted to costs. It is assumed that it is together with material just an unlimited input factor. Therefore everything is abstracted to costs and we don't differentiate between different kind of inputs. Again, this is obviously a significant simplification and therefore an ideal candidate for extensions, which introduce the concept of a real workforce, divided by skills and therefore different labor costs. Also the amount of time for production could be taken into account and with limited resources the decision for a 'suboptimal' production process which just needs less time or lower qualified workers seems a rather realistic result.

The output of a process is defined by the properties of the resulting product. These properties are described by a *vector* of values between 0 and 100. The whole model works very abstract, so we don't define specifically what these properties are and therefore also not, what these values mean. Each part of the vector just defines a range of possible states which a product could have regarding that property. For example we could imagine that our product is a car and the property is maximum speed. Then 100 would be the maximum possible speed that a car could have and 0 the least possible. However it has to be highlighted again, that the model is generic on purpose. The main function of these properties is not to find out, what kind of products develop over time. It is to have a testing mechanism, because the customers' taste is characterized by the exactly same vector. Figure 5.2 demonstrates this with an example. The gap results into a percentage value of how much the product corresponds to the customers' taste. This value then defines the demand and price at the market - both values are derived from what the ideal product would yield, but with a discount. We could imagine that a certain share of the customers interested in the ideal product still find this new product 'good enough' to buy it, but at a certain lower price.

New Product	90	80	32	65	5
Customer Taste	85	80	54	61	0
Gap	5	0	22	4	5

Corresponds 92,8% to the `ideal product'

Figure 5.2: Product vector and customer taste - The gap shows how far the product is from the ideal one, and therefore influences demand and price.

The whole process is driven by the circumstance that customers are always hungry for something new. We can imagine that for every market there exists a prototypic product that would fulfill the customers' wishes completely. What this market is exactly (if it is a whole industry, a product category or even a smaller segment) is not important, anyway this abstraction is already reducing the amount of different demands drastically, again for the sake of simplicity, but also to be able to keep an overview of what is going on in the model at all. These different demands are, as already mentioned, also described in form of a vector - it's the ideal product that the market wishes for. This taste (I am a little bit reluctant to call it demand, as we will see that the model really tries to not be too demand driven) will modify over time. The requested ideal values of the product properties change randomly over time. This leads to the development, that established products will fit this taste less and less and therefore fall in market share and/or price, as both is linked to the fitness.

So how does the combination of two processes take place? Figure 5.3 gives a schematic overview,

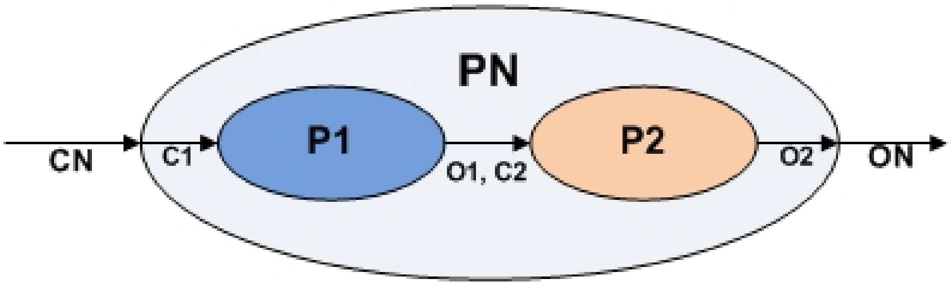

Figure 5.3: Combination of old Elements (P_1, P_2) to new process P_N

but additional explanations are necessary. Both old processes obviously had an individual cost input (C_1, C_2) and an output in vector form (O_1, O_2). The new process will have new costs C_N and a new output O_N. The first question that should come to our mind if we look at the figure, is the following: Is the link between the two old processes really directional? If it would be, then a change of direction would obviously result in a completely different product. But this doesn't seem to make sense. Innovations are rarely of the form that you just perform known processes after each other and something completely different will arise. If it would be that easy, then almost no creative energy would be needed to create novelty. It is the link itself that creates new elements, the *way* of how these old elements are combined. Therefore it seems that usually these two old processes are interwoven in a much more complex way then just pure chronology.

But even if that would be the case, then it obviously took a lot of research effort in our model to find out how to connect the two processes. For that reason alone we can not assume that this research breakthrough would allow them at the same time to figure out how to connect these elements in the other direction. But these are already way too much theoretical thoughts. We define that there is only one reasonable way to connect two elements in our model. If there is an implicit direction or not is not relevant for us and therefore also not visible in the model. However, a more interesting thought than directional changes would be if there exist more than one connection between two elements in general. It seems realistic that researchers after some time come up with a more efficient way to connect two elements which where already connected before - this new way of combining could use less resources or lead to a similar but improved product. Therefore the iterative improvement of already known connections is for sure worth to think about as an extension of the model.

The second question that arises is how these new parameters of the resulting process are determined. If we think about the costs then it seems to make sense that the new costs should be at least lower than the sum of the two old ones. While in reality maybe we would find counterexamples to this hypothesis, it seems to be plausible to assume some synergy effects. For the lower border of this range we have to make even stronger assumptions, namely we determine that the new costs can not be lower than two thirds of one of the old processes alone. If we assume for example a new combination where one process is such a revolutionary way of doing things regarding the second old process, that the costs of the latter are now negligible, then we still would have the costs from the first process. Of course we might find examples in the real world where there is so much synergy that the new costs of both process parts sink in such an drastic extent that the new costs are even less than those of either parent. Therefore we included the two thirds rule - also to have some productivity increase in terms of costs in some cases. We might sink the lower cost boarder even further if we see that new products get more and more expensive, as we don't really have inflation in our model per se.

The exact amount of the costs within that given range is determined by the link, which has this value randomly assigned. To put it in a short example: Let's assume that the first process costs 100 Euro per Unit and the second one 50 Euro per Unit. Then the costs of the newly generated process will be somewhere between 33 and 149. Of course there won't be a completely random distribution, but more likely one which is more heavily on the side of low costs, as we assume that technological progress is leading to more effective output. But this is subject to fine-tuning of the model, as certain balance mechanism between the costs on the one hand and the price they can charge for the resulting product on the other hand (which is determined by the market) will have to be implemented, otherwise the companies will get all super-rich or die out very quickly. While this seems rather generic and done for the sake of the model, one also has to be reminded that the aim of this model is not at all to simulate a whole economy. Therefore naturally some connections and constraints are missing. The market is only implemented as a feedback mechanism and to 'steer'' in which way the developments are going, but it is not a simulation of real consumers.

Speaking of the market feedback: The probably more interesting 'fusion' is taking place in terms of the final output after the combination of two processes. Our assumption is, that in general there will be a tendency to hit the middle ground. That means if for example process 1 resulted in a value of 80 for a specific property, and process two had a value of 60, then the final result will tend to 70. This is not only the simplest approach, but it also satisfies the need for a certain predictability. While the producers can have different approaches for which combination they chose to research we still can assume that the expected resulting product has to play a huge factor. A certain randomness nevertheless has to exist obviously. Otherwise we would have complete forecast and that would not only kill the fun, but also the purpose of this model. Therefore we will allow some kind of random drift in the direction of either of the two old processes. So in our example from before it is very likely that instead of 70 the result could be for example 68 or 73. As such a mechanism alone would lead in the long run to a lot of outputs in the middle range, we also introduce a certain mutation chance, where single properties just gain completely random values again.

Another conceivable idea is to introduce the concept of 'supportive' processes, where other processes are only supported by them, and therefore the properties will always tend stronger into the direction of the dominant one. This might sound strange, but if we take a very simplistic example and we think that one process is the general idea of creating a table and the other one is how to cut wood, then we could assume that the result would still be very similar, even if we cut the wood with machines instead of manually. These supportive processes then would mainly increase the productivity of others (which means reducing costs in our model) but not alter the end-product significantly. But for now this is only a possible extension, as on the one hand it would enable a nice way to introduce real progress (we will discuss very soon, that this is kind of a problem) but on the other hand we would have to think that our agents are extremely stupid if we assume that they don't find out, that mixing other processes again and again with a supportive one would give continual productivity growth. One way to do it would be to keep track and make combinations which have already included this supportive process in their history not feasible to combine with that process again. But this would need a lot of additional 'memory' in each process.

To come a little bit more back to our basic 'procedure': When the research of a combination is finished and it turns out that a combination is feasible, it has to go into the market test. Obviously the producers don't know the exact taste of the customers, otherwise such a test would be obsolete. The market test returns an approximate value of the fitness, the possible price and the market size and share - however all of these can change anytime, may it be because of changing taste or new competitors. It is also planned to introduce the concept of trend setters. As already mentioned earlier, the model shouldn't be too demand-driven. Therefore every new product that is tested has the potential to create a whole new demand, a whole new market as we called it earlier. The more distinct it is from all other currently requested products, the higher the probability that it 'fills a void of which nobody knew before that it even existed'. We then assume that the 'ideal' requested product of this new introduced market is still quite far away from the product which created this demand. Otherwise we wouldn't have much dynamics going on.

Actually this is even rather easy to explain: The new product gives a certain glimpse for what is possible in that direction, but of course it is still far from perfection. For example when the first personal computers where introduced, this created a whole new market, but still the first products obviously lacked a lot of features needed. As we have now a mechanism to create new demands, the question is if we also need one to eliminate old ones, as otherwise the number would constantly grow. As this might be even realistic, we could still imagine that some demands - even though they are updated with mutations of the requested properties all the time - die out after some time.

There exist obviously different possibilities on how to implement this. One idea would be to just erase them with a certain, very low probability. Another, even more tempting one would be to just reduce the market size of this demand over time so that it kind of fades out. The way it actually is implemented is similar in the basic idea, but executed differently: As our markets are anyway changing their taste all the time it happens that after a certain amount of time they land in the 'extreme zone' - that means, their requested properties are all either extremely high or low. Such an extreme market could have a taste vector like 89|95|93|91|87. In our implementation we will eliminate these extreme markets as we interpret their development as having become irrelevant.

We are almost done now with the topic of the market feedback and the emergence of new combinations and products. However, as already hinted earlier, one big question mark remains unsolved: How to model real, time-irreversible progress. Of course this model provides permanent change and a dynamic development, but there is no real indicator that what is produced in time 1000 is more advanced than in time 1. In both time periods the customers could per accident demand exactly the same product, as the properties of our products are not indicated in the way that one end of the range (100) is very good and the other end (1) is very bad. They are just properties, without any quality indication. This could even be justified, as especially with fashion there is the saying, that everything is coming back again after a certain amount of time. But in our model our producers then could even take the step back and just revive the old, long abandoned processes to produce these products, which seems rather unrealistic. Therefore the model has to ensure in some way, that new processes are more relevant than their forefathers. There are several possible ways to ensure that:

1. The restrictions regarding the distribution of possible costs of the new process have to be even stronger in the direction that new processes are never costing more than their ancestors.

2. The concept of iterative upgrades of already existing connections would ensure that even if they are re-using old combinations, they are trying to do it in a new and more productive way.

3. The concept of supportive processes would be an integrated mechanism to ensure that new processes which are a result of combinations with these newer supportive technologies are always more productive than old ones. As we already saw, this would bring some implementation problems as every process would have to have a perfect memory of all its ancestors.

4. Old processes just vanish after a certain amount of time. This is not because they are forgotten but just because they wouldn't be productive anymore.

For the sake of simplicity we choose the last alternative, but it is obvious that especially the second and third would add even more 'spice' to the model.

Production & Competition

If a market test was satisfying enough, a producer can decide to give the new product a 'Go'. The decision to produce or not is for sure not always easy, and depends mainly on which other products it has and if it has better options to spend its limited budget or not. The production as well as the research is costing money, which is our abstraction of input of workforce and material. In our model it could happen that a company decides to produce something which is far away form the ideal product that the customer wants, but it hasn't anything else to produce which is more profitable and/or the competition on the market is so weak that even a bad product is selling good. Or it just discovered a very efficient way of 'mass production' so that quantity goes over quality (as long as the demand on the market is there). It might even prefer this cheap mass production over another process of which it has knowledge which would result in a better product, but just doesn't pay off in terms of costs and chargeable price.

As already implied all companies have a budget on which they are running on, which is at the same time their capital stock. There will be some indicator of the company size (depending on the amount of production in the past) and based on that some fixed costs. Again, the aim is not to simulate a realistic economy. However, this fixed costs are necessary, otherwise a company could prevent it's own death just by stopping to spend money on research or production. There has to be pressure for them to do constant research. If a company is running out of money, it will die out. Also, the emergence of startups is possible (usually with a high specialization in one new product). In the case of our model this happens if a market surpasses a certain growth rate and with a certain random chance. The interpretation to this is that fast growing markets are most attractive to found new companies with a highly specialized product for that market.

So, as we see, in our model markets as well as companies can die and emerge dynamically. The last question to be answered is, when are companies researching and when not? Obviously there exist a lot of theories on this topic. Based on the possible developments of this model it seemed to make sense to implement three different conditions under which a company is investing into research. It has to be admitted that despite the circumstance that they can reasonably explained by looking at the real world, they are also chosen to keep the simulation runs as dynamic as possible. Obviously once again this topic is a good candidate for extensions. The agent-based pattern of the model would make it easy to find different types of company agents with different innovation strategies. However, the three conditions of this basic model are the following:

1. The company has no profitable product on any market. As the fixed costs would drag this company down to its extinction anyway, it has to research in a desperate attempt to turn things around. With the additional research costs its capital will decline even faster, but at least there is a chance to have a discovery quick enough to survive.
2. The company is doing great and has the freedom to research to stay in this good position without too much financial worries. Technically, a 'top company' is indicated by being in the X number of companies with the highest capital and having Y years of growth.
3. The company still receives public innovation funding. Every start-up in our model gets a funding for research for the first couple of defined years. This funding is decreasing and then stopping totally at some point in time. As the funding is never below 50% of the total research costs, companies who are eligible for funding will always use it.

For each market there is a defined market size. The market share and the price for each product is defined by its closeness to the customers' ideal product, but also of course by the actual production volume of a company. If a company produces below it's maximum possible share (because it doesn't have enough budget) then probably others will fill the void. Of course we know from reality and standard economic models that higher demand than supply usually rises the price, but in our case it would have gotten out of hands and would make the model even more complicated than it already is. The production itself just simply happens by paying the costs for each piece and then it will be sold on the market immediately, if there is the corresponding demand. In this basic version of the model companies do not have the possibility to store their items, so every overproduction is immediately useless. Obviously the patterns in term of competition are expected to be one of the most interesting observations of this model - Will there be a constant dynamic of new fast climbers, will it be a balanced competition or will - as it is the case in reality - a few big companies dominate?

5.2 Agent-based Implementation with NetLogo

In the previous section I already tried to explain the model as far as possible in a rather non-technical fashion. Now it is time to look at the actual implementation which was done in NetLogo. When we look at figure 5.4 then it can be quite overwhelming. Especially the huge amount of sliders, where many of them have quite a big influence on the model (we will come to this analysis later), can be confusing. Therefore I will go over all of the user interface elements and will try to explain them in some sort of technical user manual.

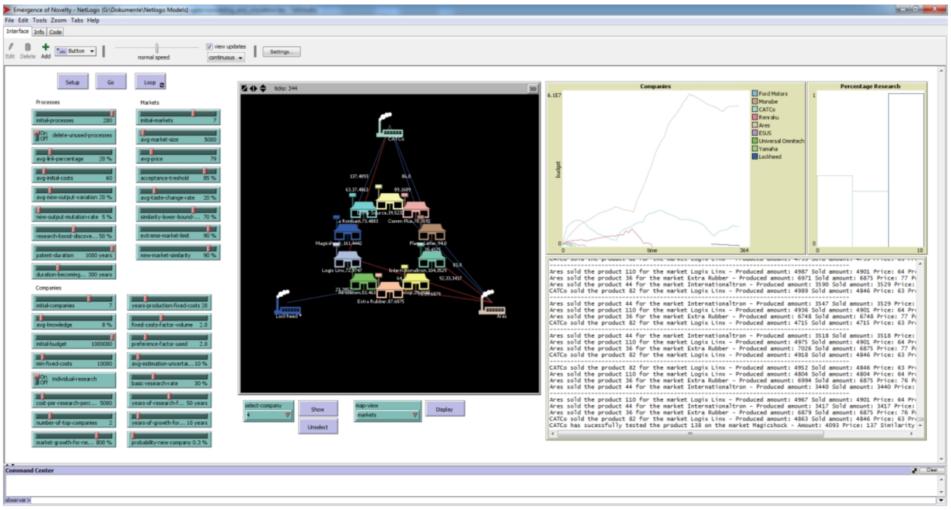

Figure 5.4: The whole Interface of the NetLogo simulation - At the first glance it can be quite overwhelming.

Before we look at the model-specific elements, the most important controls are the three buttons at the top (see figure 5.5), which are rather NetLogo standard. 'Setup' has to be pressed after all parameters are chosen to setup the simulation with these parameter values. The 'Go' button lets the simulation run for exactly one *tick* - one time unit (in our case the equivalent of a year) so to say, in which all methods are executed once in the order defined by the coding. 'Loop' on the other hand does nothing else than also calling the 'Go' method, just for infinite times until the button is pressed again to stop it. This is the choice you usually prefer when you just wanna see the simulation running and observe the dynamics.

Figure 5.5: The three main buttons for running the simulation

The World View - Companies, Markets and the Process Network

Let's start with the flashiest element, which obviously draws the eye, as there is most of the visible dynamics going on: The world view. Experienced NetLogo users will notice that our implementation is not using the common turtle-patch paradigm - at least not in the usual way. Our agents are technically very well *turtles*, but there is no movement over the world. Instead our focus lies on the connections (called *links* in NetLogo) between our agents - implemented as custom link breeds in NetLogo, and these are changing rapidly.

The first thing that has to be noticed, is that there are two world views available: The markets view and the processes view. In figure 5.6 you can see the markets view. Before we go into the details of this view, also the control elements below the view have to be briefly explained. On the right side we have the choice between both mentioned views. If the simulation is stopped, then a click on the display button is necessary to update the choice, otherwise the next tick will have the same effect. With the chooser on the left you can select which company should be displayed, but this is only relevant for the processes view and will therefore be explained a little bit later.

For now let's have a look at the markets view. What we can see here is that all markets are placed in the inner circle, represented by some kind of store building. The companies are arranged in the outer circle and their graphical representation is a factory building. But what else does this view tell us? First of all we can see that markets as well as companies have names, those are randomly assigned from a given list and should help in following the events of the simulation easier. In the case of the companies there is also a '*' next to their name in the case that this company is investing in research at that moment. Markets on the other hand have two numbers next to their name: The first one is the maximum price achievable in this market (with the 'perfect' product) and the other one is the potential market size, which means in our case the maximum number of products that can be sold in this market per year.

Nevertheless actually the most important information lies within the connections between markets and companies that we see. A blue line indicates that the company would have in general a fitting product for this market (where fitting means it surpasses the given similarity threshold). It is a market estimation and shows the estimated price for the product as well as the estimated amount of products that could be sold. It has to be noticed that these estimations are not totally correct, as first of all there is an estimation error and secondly the estimation is based on last

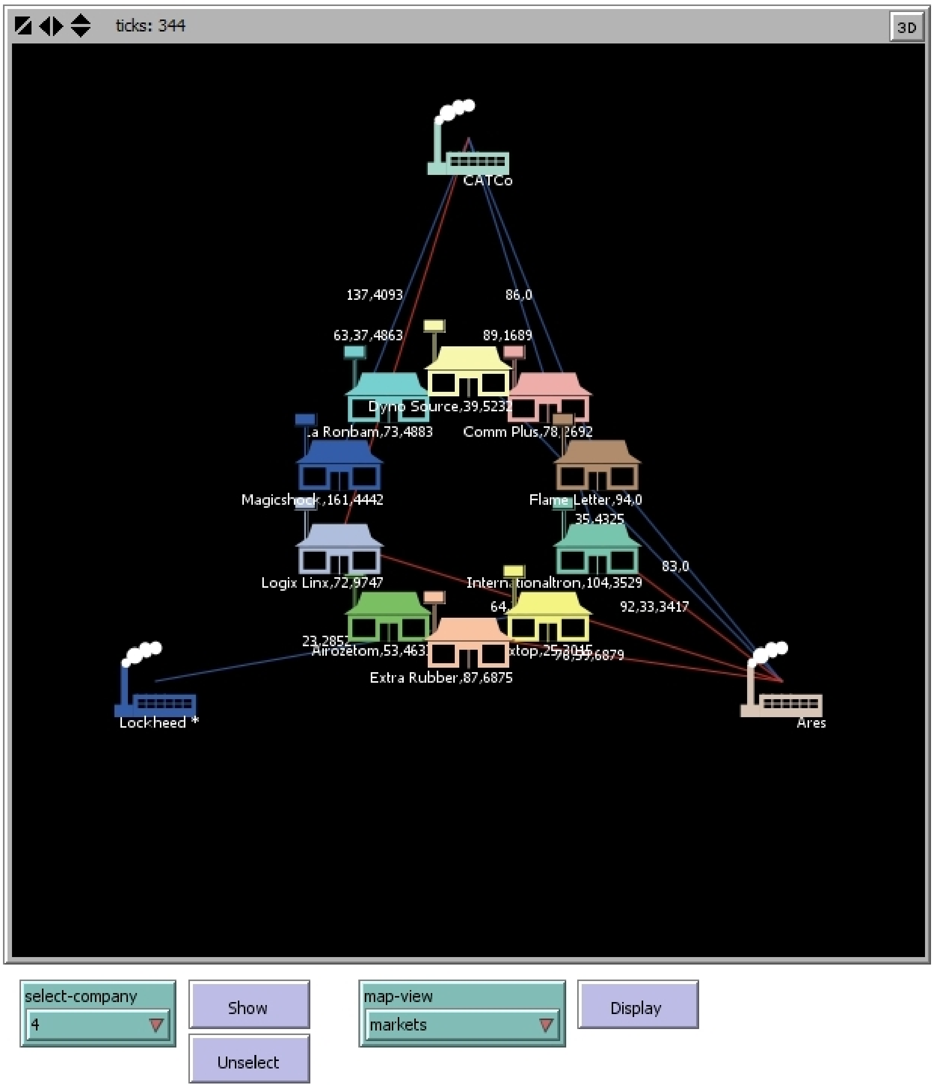

Figure 5.6: The world view with the markets view displayed

years competition, therefore can not take into account competitors which started to produce for this market in the same period.

As long as we see a blue line, this means that the company is not actually producing yet for this market. This can have different reasons. One could be that the product was just recently researched, so it is only in the next period possible to start producing and selling it. Another one could be that given the production costs this product just wouldn't be profitable. A third common cause is that the company is already producing other products for other markets and just doesn't have enough budget to produce this product as well - this obviously implies that the other products are more profitable as we don't have any costs or delays for changes of production, which therefore would be always possible immediately from one year to the next one.

As soon as a company is actually producing (and usually also selling, just maybe not with the expected amount) for a market, the blue estimation connection is replaced by a red sales connection. Now we have three figures instead of the previous two. The first one is the actual price that customers will pay, the second one the revenue per piece (price - production costs) and the third one is the produced amount. Please note that the sold amount most likely will be different.

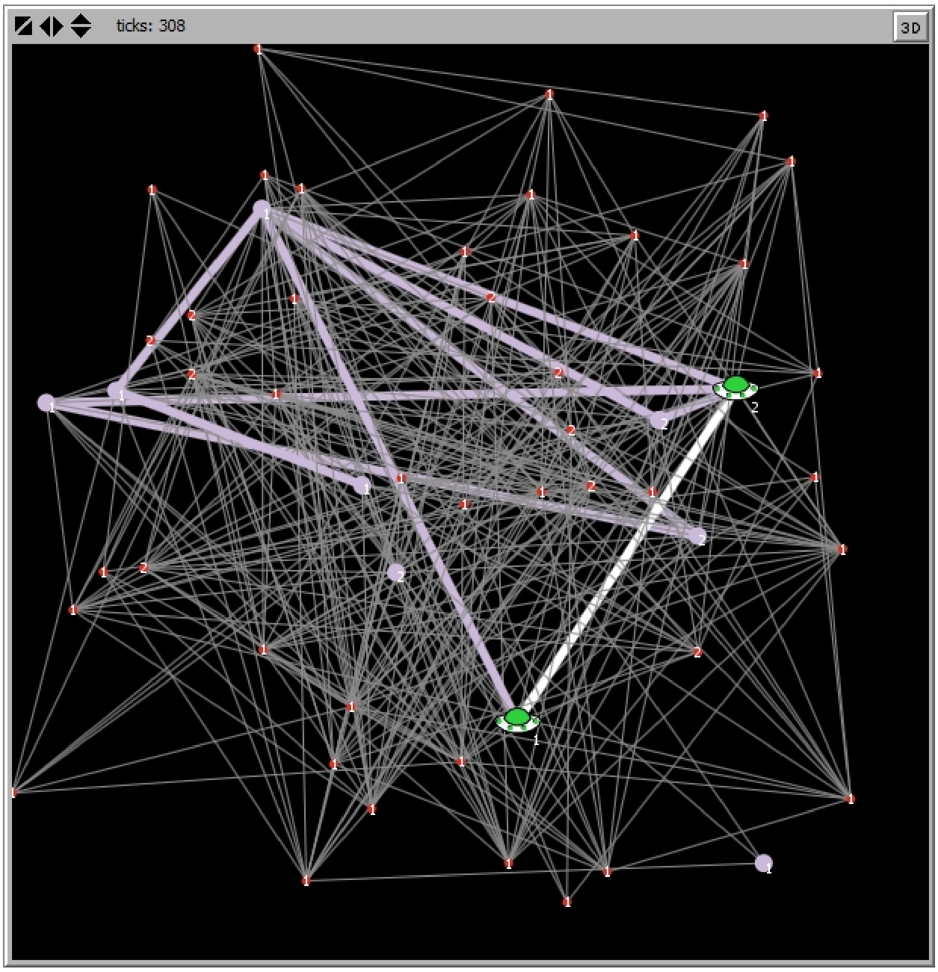

Figure 5.7: The processes view - The selected company is highlighted

It is of course not possible to display all relevant information this way, especially as it can depend from the situation or observer what is considered relevant. However, to address this issue there are also other means of output to give a more complete picture, which we will discuss next. One thing that shouldn't be left unmentioned is the 'inspect' functionality which is included in

NetLogo. With a right-click on any element (be it turtles or links) in the view you can choose to inspect it and then you get a detailed list of all its properties and variables.

The second view that is available is the processes view (see figure 5.7). Here you can see the network of all researched (and not yet irrelevant) processes and their connections (as far as they exist). As this can be quite a big amount of nodes, especially when there are several companies alive, there is the possibility to select a single company. The processes and their connections which belong to the knowledge of this selected company are then highlighted in the color of the company, which is also used for the other view and plots.

As it is a standard function, here the 'inspect' functionality can again be of big use to inspect processes as well as their links to see costs, the resulting product vector, and so on. For a better picture of what is happening in the simulation the two processes which are researched by this company are highlighted even further (with the UFO symbol, which can also be seen in the figure). If there is a connection between those two processes, which means that the research will result in a new process/product, this connection is highlighted white. The last information this view shows us is a number next to each process. This indicates how many companies have knowledge about said process.

Other Outputs - Plots and the Event Log

To get a clearer picture on what is going on in the simulation run there are two plots and one event log output provided. The first plot might be considered the most important, as it shows the development of each company in terms of budget over time, and therefore gives an overview of the rises and falls of these competitors. New companies are automatically added with a new plot pen at the time mark where they joined the 'race'. So the development over time is perfectly traceable. An example of both plots can be found in figure 5.8.

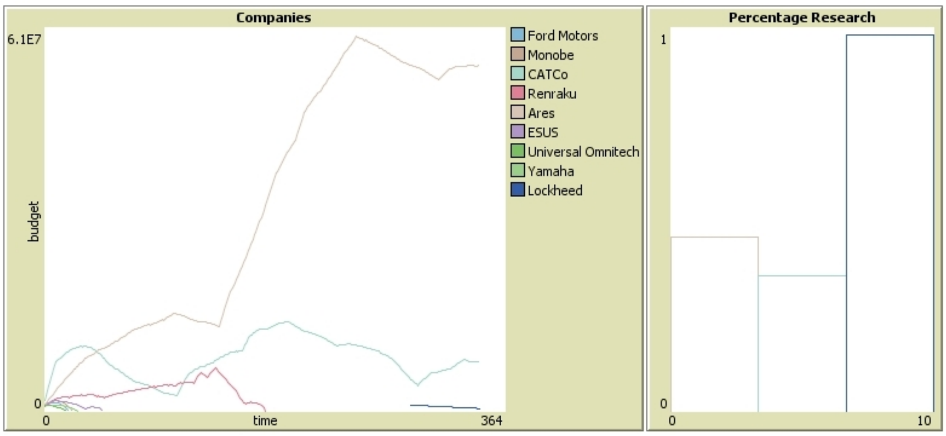

Figure 5.8: The two plots: On the left the performance (budget) of the companies over time, on the right the research rate

The second plot could seem a little bit unspectacular, but it bears also interesting information: It shows the percentage of research years compared to the total of years the company lived of all currently active companies in a bar graph. If for example a company exists for 500 years but has only in 50 years of this timespan decided to research, then the research percentage would be 10%. This should give an overview how companies with different research rates are performing in comparison to each other, or to put it differently: To see if there is any connection of being among the leaders and a high research rate.

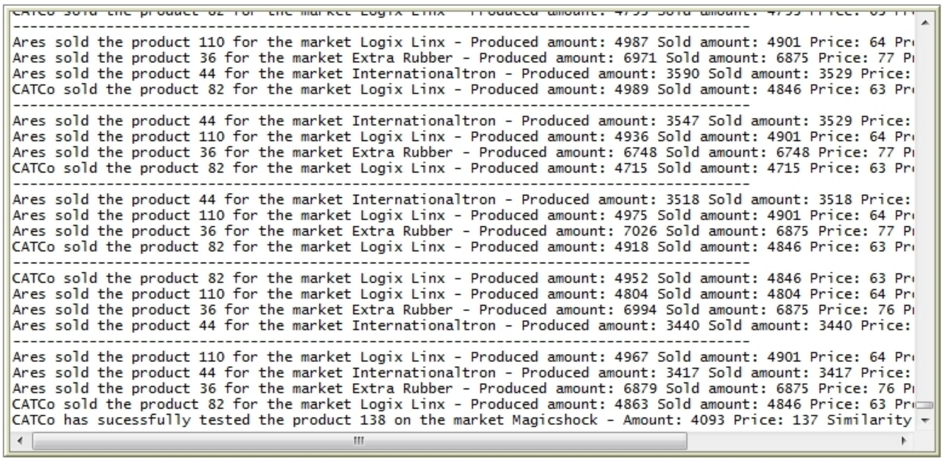

```
------------------------------------------------------------------------------------
Ares sold the product 110 for the market Logix Linx - Produced amount: 4987 Sold amount: 4901 Price: 64 Pr
Ares sold the product 36 for the market Extra Rubber - Produced amount: 6971 Sold amount: 6875 Price: 77 P
Ares sold the product 44 for the market Internationaltron - Produced amount: 3590 Sold amount: 3529 Price:
CATCo sold the product 82 for the market Logix Linx - Produced amount: 4989 Sold amount: 4846 Price: 63 Pr
------------------------------------------------------------------------------------
Ares sold the product 44 for the market Internationaltron - Produced amount: 3547 Sold amount: 3529 Price:
Ares sold the product 110 for the market Logix Linx - Produced amount: 4936 Sold amount: 4901 Price: 64 Pr
Ares sold the product 36 for the market Extra Rubber - Produced amount: 6748 Sold amount: 6748 Price: 77 P
CATCo sold the product 82 for the market Logix Linx - Produced amount: 4715 Sold amount: 4715 Price: 63 Pr
------------------------------------------------------------------------------------
Ares sold the product 44 for the market Internationaltron - Produced amount: 3518 Sold amount: 3518 Price:
Ares sold the product 110 for the market Logix Linx - Produced amount: 4975 Sold amount: 4901 Price: 64 Pr
Ares sold the product 36 for the market Extra Rubber - Produced amount: 7026 Sold amount: 6875 Price: 77 P
CATCo sold the product 82 for the market Logix Linx - Produced amount: 4918 Sold amount: 4846 Price: 63 Pr
------------------------------------------------------------------------------------
CATCo sold the product 82 for the market Logix Linx - Produced amount: 4952 Sold amount: 4846 Price: 63 Pr
Ares sold the product 110 for the market Logix Linx - Produced amount: 4804 Sold amount: 4804 Price: 64 Pr
Ares sold the product 36 for the market Extra Rubber - Produced amount: 6994 Sold amount: 6875 Price: 76 P
Ares sold the product 44 for the market Internationaltron - Produced amount: 3440 Sold amount: 3440 Price:
------------------------------------------------------------------------------------
Ares sold the product 110 for the market Logix Linx - Produced amount: 4967 Sold amount: 4901 Price: 64 Pr
Ares sold the product 44 for the market Internationaltron - Produced amount: 3417 Sold amount: 3417 Price:
Ares sold the product 36 for the market Extra Rubber - Produced amount: 6879 Sold amount: 6875 Price: 76 P
CATCo sold the product 82 for the market Logix Linx - Produced amount: 4863 Sold amount: 4846 Price: 63 Pr
CATCo has sucessfully tested the product 138 on the market Magicshock - Amount: 4093 Price: 137 Similarity
```

Figure 5.9: The event log shows what is going on in detail

With all the already mentioned output elements the observer should get a pretty good overview of the general progress of companies and markets. However, even with the ‘inspect’ function it is not really possible to understand what is happening in detail in every single tick. Therefore we also have the event log (see figure 5.9). Here all the most important events are printed in text form. It tells us for example when a company successfully researched a new product, tested it on the market or started selling it. Other important events that can be found there would be the emergence and death of companies and markets. Of course, while running in ‘Loop’, this log - even though it was tried to reduce it to the most important events - will provide a vast flood of information. But this is not what it is intended for. For single ticks it is indeed very useful to get an idea why the changes that we can see in the markets view actually really happened. Also it gives a good overview of all market transactions in that period - for example how many exemplars were sold compared to the amount produced by a single company.

Control Elements - Sliders and Switches for External Parameters

There is a huge amount of external parameters in the model, and while some might be hardcoded, most of them can be changed by the user. But this also leads to a complicated looking user interface. Therefore in the implementation these parameters are grouped into three categories, which are the agents to which they should fit best - processes, markets and companies. Obviously some would belong into at least two categories but it was tried to group parameters which are connected to each other.

In the following three tables the respective groups are presented and you will find a short description of each parameter including a discussion of its recommended values.

Processes

Parameter	
Description	**Values and Recommendation**
initial-processes	
This is the number of processes which are generated initially. Based on the avg-knowledge of the companies each company then receives its share. All processes which are not connected to a company are deleted (if delete-unused-processes is set). Therefore even after pressing the 'Setup' button in many cases there will exist less processes than selected here.	The higher the selected number, the less probable it is that a process is shared by several companies. As diversity of knowledge is in general good in this model, it is recommended to choose the maximum value of 200.
delete-unused-processes	
This switch decides if after the initial processes are assigned to companies, the unused processes should be deleted or not.	For now the clear recommendation is to set it to 'On'. This parameter is only necessary for a possible extension where processes can be added to the knowledge which are not the result of a combination.
avg-link-percentage	
It defines with how many other processes each process is connected via a link, given as a percentage of the whole set of processes.	There is not really any recommendation, although the value has obviously some influence on how often a research will be successful. A good starting point would be around 20%. Higher values can also affect the performance as the number of links is then growing enormously.
avg-initial-costs	
This defines the average costs of the initial processes (gamma distribution) and therefore also affects the general costs level of resulting combinations.	The level of the costs value itself is not as important as its proportion regarding the other cost and income parameters - to avoid both, too frequent death and too easy growth. Therefore throughout this tables we will just give one overall example setting which delivers quite good results. For avg-initial-costs this would be a value around 60.
avg-new-output-variation	

Defines the average deviation of a ‘child’s’ product vector from the average of both ‘parents’.	Depending on how much predictability you want this value should be low for a high predictability and vice versa. Right now it doesn’t have much influence as companies are not really expecting anything, but can be important for extensions in that regard. 20% to 25% as a starting point seems to make sense.
new-output-mutation-rate	
Besides the variation there is also a chance for total mutation of a child’s vector property. This is defined with this parameter.	Again it depends how much predictability you want. But this parameter is also important for the purpose that in the long run not all processes end up in the middle range of possible values (which is of course the case when taking only the average all the time). A good value is around 5% to 10%.
research-boost-discovered	
This parameter defines how much faster a company can research a process which is already discovered by at least one competitor. For every additional company that has it already discovered the boost is multiplied with itself, which means that it’s usually getting less with every company. If the boost would be for example 50%, it would be 25% for the second (so together 75%) and so on.	A value around 50% seems to be right with an implementation like that. This way the research boost will never be higher than doubling the original research speed. Too high values lead to an automatic research success after one period.
patent-duration	
This defines the time span until when a process becomes common knowledge to everyone (beginning with its creation).	For our results a rather high value, usually even the maximum of 1000 has created the most interesting results, which means that the whole concept of becoming common knowledge does not work too good with the current implementation - companies become too fast too similar to each other.
duration-becoming-irrelevant	

With this parameter we can define after which time span a process which is not used (that means it is neither sold as a product nor used for research) gets deleted.	It seems like a value in the middle ground like 300 to 500 leads to good results. This is on the one hand the engine for supporting progress and therefore rewarding research, but it also has performance purposes (without deleting the old processes it would get slower and slower). A too low value results in a situation where companies are losing their knowledge much faster than they can generate new one.

Table 5.1: The external parameters related to processes

Markets

Parameter	
Description	**Values and Recommendation**
initial-markets	
Simply defines the number of markets at the setup phase.	There is no recommendation - fewer markets just mean that there will be less diversity of successful processes - which usually also means that less diversity among companies exist, and probably fewer will survive.
avg-market-size	
This parameter sets how many units at average can be sold at best on this market (the companies have to share this amount). The actual market sizes are then determined randomly by a normal distribution.	Like most of the prize values this parameter is more about the ratio to other parameters than the absolute value. For our example we chose 7500. A bigger market size means that more companies can survive by selling to the same market, putting less pressure on research.
avg-price	
Defines the average price for which a product can be sold. Again the actual maximum price of the market is then defined randomly by a normal distribution.	Again, it is more about the ratio to other parameters than the absolute value. For our example we chose 80. A lower value (especially regarding the production costs) leads to more unprofitable markets. Too high values make it too easy for companies to earn money on one single market.
acceptance-threshold	

This is the threshold in terms of which similarity is accepted for a product to be sellable on a market. Only products with a similarity which is at least equal to the threshold can be sold there. As market tastes change, it can happen that a process is falling below this limit from one year to another.	The higher the value, the less products are useful in the end. What is even more important is that with a threshold not too high, it is more likely that different companies sell different products to the same market - which is what we would expect based on reality. However a value too low would mean that almost every product could be sold on every market. The middle ground is therefore somewhere between 80% and 90%.
avg-taste-change-rate	
With this setting we can define how often on average a market taste property changes, whereas a taste change means +1 or -1 of the vector element, depending on the current taste change trend, which can be positive or negative. A value of 50% would then mean that every vector property is changing on average every second year.	A rate too high makes it impossible for companies to follow up with new researches and hardly makes any market profitable in the long run. A value too small and there is no incentive to research at all. Our experiences recommend values around the 10% to 25% range.
similarity-lower-bound-new-market	
This parameter is relevant for the creation of new markets. When a company tests a new product on the market and it is in terms of similarity to every existing market below or equal that limit, then the product creates its own market, because it is so innovative.	The way similarity is calculated, values below 50% are almost impossible to reach. Therefore the middle ground should be in the 70% area. Values too low don't create new markets, values too high let them sprout too easily.
extreme-market-limit	
Here we have the counterpart to the previous parameter, as this one defines when a market is going to vanish. With the taste change trends, markets tend to take on rather extreme vector values, very close to either 0 or 99. If a market reaches the here-defined similarity to a fictive product on one of the both extremes (0|0|0|0|0) or (99|99|99|99|99), it gets deleted.	Obviously a rather high value is needed to make this a rare occasion (which it should be). Therefore it shouldn't be below 85%, better even 90% and upwards. This depends a lot on the value of similarity-lower-bound-new-market. These two parameters should lead to an almost equal speed of birth and death - one good working example already given is the combination 70% / 90%.
new-market-similarity	

This simply defines the similarity of a new market to the product which created it. The thought behind this is, that rather seldom a new product immediately fulfills the customers needs perfectly, but usually is only a first step in that direction and gets refined later.	There is no real recommendation with this one, as it mainly depends on how much of an advantage we want to give the company which came up with the product. The only restriction that seems to make sense is to keep the value above the acceptance-threshold, otherwise the product can't get sold on the market it created. So for an acceptance-threshold of 85% a new-market-similarity of 90% would be feasible.

Table 5.2: The external parameters related to markets

Companies

Parameter	
Description	**Values and Recommendation**
initial-companies	
Simply defines the number of companies at the setup phase.	There is no recommendation, but fewer companies mean less competition, but also less chance that any company will survive the existence struggle in the beginning, as also the pool of processes will be much smaller.
avg-knowledge	
This defines the share of processes of the total pool that a company can add to its knowledge at the time of its creation.	Again it is hard to recommend any specific value here. As it is creating more interesting results when the knowledge is rather fragmented, this value should be kept rather low, probably around the 8% to 10% mark. Higher values also can become a performance issue, as they mean that the simulation starts with a bigger process pool (as all unused processes are deleted).
initial-budget	
As the name says, here we can set the initial budget of the companies.	In the long run this parameter is not really important. However a higher value can help that in the starting phase more companies survive, which usually leads to more interesting runs. Of course it could be that this is not desirable, then a lower value should fit better.
min-fixed-costs	

The main part of the fixed costs is calculated based on the company size, which is derived from previous production volumes. However as it would be theoretically possible that a company could survive by simply not producing and therefore not losing any money, this parameter was introduced and simply subtracts a fixed amount every year, so that production and/or research is necessary to survive.	As this parameter is really only there for making inactive companies die, it shouldn't be too high, as otherwise bad economic periods get punished too hard and no company will survive very long. In the long run even low values like 5.000 to 10.000 hurt enough. However it also depends on the general cost/price level.
years-production-fixed-costs	
This parameter defines how many years of past production volumes are counted for the determination of the size of the company. For example a value of 20 adds up the production volume of the last 20 years and multiplies it then with the fixed-costs-factor-volume to determine the fixed costs.	The recommendation for this parameter depends on how we want to treat big companies as soon as their rise begins to stop. With a high value they will fall very fast back to the ground without any new sales, while a low value will allow them to float on the top for a very long time. Of course it also depends on the combination with the fixed-costs-factor-volume. A high value for the years and a low value for the factor can lead to the same total costs like vice versa, but in the first case the company is 'haunted' longer by its former success but falling down slower, while in the second case it would pay a lot in the first few years of misfortune, but would finish its downsizing much faster.
fixed-costs-factor-volume	
This setting is already explained with the description of the previous one. The years-production-fixed-costs are multiplied with this factor.	See years-production-fixed-costs.
preference-factor-used	

In our current implementation the decision which processes should be researched are based on random choice, only those processes which already resulted in a successful combination are preferred. This parameter defines the factor of preference - for example a value of 2 would make it doubles as likely that already used processes are chosen again.	A value around 2 seems plausible. Too high values would lead to an over-usage of preferred processes.
avg-estimation-uncertainty	
With this parameter we can set how much a company is on average wrong when making a market test and estimating the future demand amount and price. The actual estimation-uncertainty in each case is then defined randomly by a normal distribution.	In general this parameter is not too important, as companies anyway get feedback after the first time that they sold a product, so even if their estimation is so wrong that contrary to their belief a product is not profitable, they will realize that after one period and can stop production immediately. However if we would have a delay mechanism and they would be stuck with a production for some time, this would have much more influence. Nevertheless the value shouldn't be too high as it seems rather unrealistic then. Something in the range of 10% to 20% seems to be fine.
individual-research	
This switch defines if companies individually 'decide' to research or not. If the switch is off, in the current implementation they always research.	To keep it short: Obviously the individual research approach is much more sophisticated and can bring more interesting results. For a definition on when companies are researching in this case, please read the previous section or look at the following parameters, as they are all related to this topic.
basic-research-rate	
Here we can define the basic success rate of research. If we choose for example a value of 50% this would mean that on average it takes two years for finishing the research.	The optimal choice for this parameter depends too strong on a lot of other settings (for example duration-becoming-irrelevant), to give any recommendation. Values from 20% to 50% or even more can all make sense - a good indicator for increasing the value is if companies are constantly running out of knowledge. If individual-research is turned off, the value can usually be lower.

cost-per-research-percentage	
With this parameter we can control how much each percentage of research costs a company each year. Of course this would be even more interesting if the company could choose the research intensity, but in the current implementation there is only the decision to research at the basic-research-rate or not at all.	As previously mentioned all the cost, income and budget parameters are highly dependent on each other and their ratio is more important than the absolute value. In our example we use costs of 5.000.
years-of-research-funding	
This defines for how many years a company receives research funding after it started to exist. As long as it receives funding, it automatically researches, even though the funding decreases over time.	While a too high value can benefit the overall process network, it can hurt a single company by stretching the time of 'mandatory' research too much, especially if the circumstances would make this company to use the money otherwise without this obligation. A value around 50 years brought good results.
number-of-top-companies	
Under the current implementation one of the conditions for research is that the company 'is doing well'. One of the criteria to qualify for this statement, is to be among the top companies in terms of budget. This parameter sets how many companies are considered to be on the top.	As in the most runs the number of concurrent companies will be rarely over 10, the parameter should not be chosen too high, 2 or 3 seems most plausible.
years-of-growth-for-research	
The second criterion for being considered a top company is to have a certain amount of years of consecutive growth, which is determined via this parameter.	If the value is chosen too high, there will be almost no research among healthy companies. It probably shouldn't be higher than 10 to 20 years.
market-growth-for-new-company	
If a market grows exceptionally fast, it can attract new startups which join the market with a new innovative product. This parameter defines how much growth there has to be, that there is a chance for creating a new company at all.	Our runs showed that this value should be surprisingly high, otherwise new companies will sprout way too often. It also depends on the next parameter, but values above 800% are delivering indeed good results, which means approximately balancing out the amount of companies which are vanishing.
probability-new-company	

This parameter defines how likely it is, that a new company emerges, given that the market-growth-for-new-company threshold is surpassed.	Again, we have to recommend rather extreme values, this time in the lower area. It shouldn't exceed 0.3% to 0.5%, otherwise the simulation will be flooded with new companies.

Table 5.3: The external parameters related to companies

5.3 Simulation Runs and Exploration of the Behavior Space

After this rather exhausting list it is time to look at the simulation runs and the influence of some of the most important parameters. I will not try any mathematical analysis approach and I will also not cover all parameters - there are just too many and the influence of them is not always expressible in simple figures, but rather they shape the whole network and interactions. Therefore I will instead pick a sample parameter constellation and will present the typical results. Also I will highlight a couple of parameters which showed especially much influence on the simulation behavior and will explore how the results are changing with the parameters.

Let's start with our default constellation. This value set is what I used most of the time during the implementation to try out new features - obviously the implementation itself and the results which I got while trying out influenced the parameter choice in a feedback loop. The specific settings you can see at figure 5.10, but they also should match all the recommended values given before.

Typical patterns

When we look at figure 5.11 we can see a lot of the typical behavioral elements that we see in a lot of simulation runs - especially with parameters in the recommended range. In the beginning for example we see that two companies clearly win the race, the others often hardly survive or even die rather quickly. Most of the times it's one to three companies which can really rise in the beginning. After some time of growth new companies join the race, but usually most of them fail. At the same time one of the two market leaders can win some ground to it's competitor. These events are most probably not connected to each other but each of them is very common.

We also see, that new companies are mostly created in batches - one shortly after another. This phenomena and it's economic relevance and foundation in reality we will analyze more in detail in the next section - together with some of the other more interesting effects. As already mentioned, most new companies don't survive very long. But from time to time there are some which can really set a foot into a market and soar. In our example this is happening at the same time as the two leading companies are losing ground. These two events are more likely to be connected than the previously mentioned, but we have also a lot of cases where the rise of a new company is not immediately affecting the current leaders.

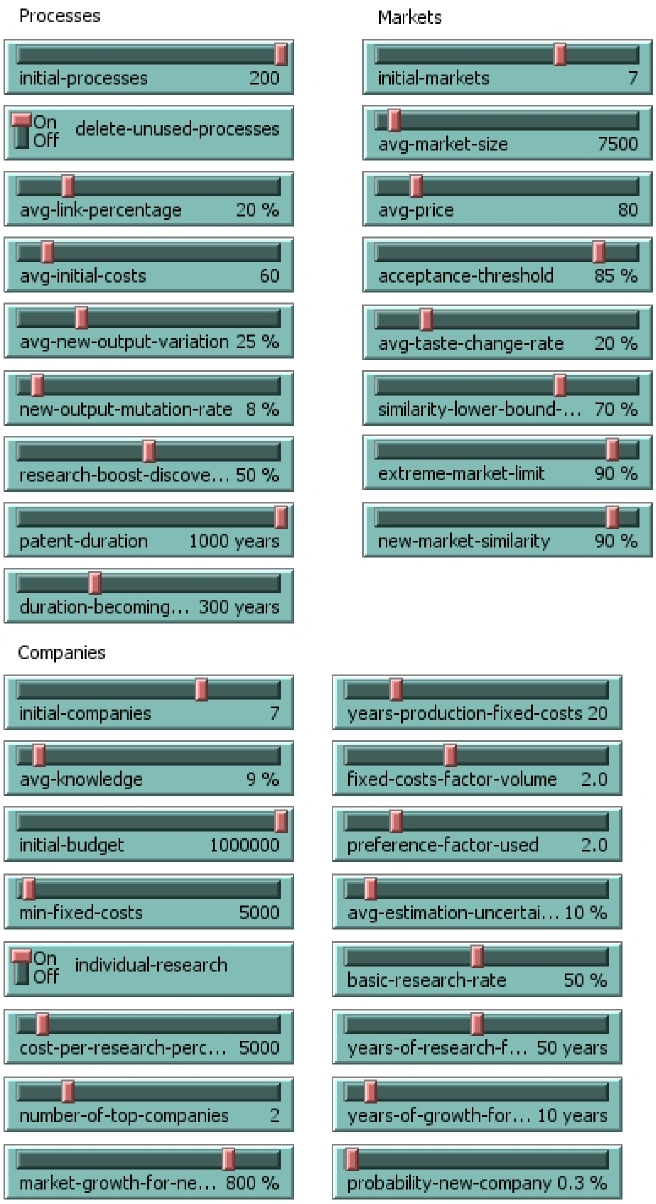

Figure 5.10: The parameter values for our ‘default’ setting

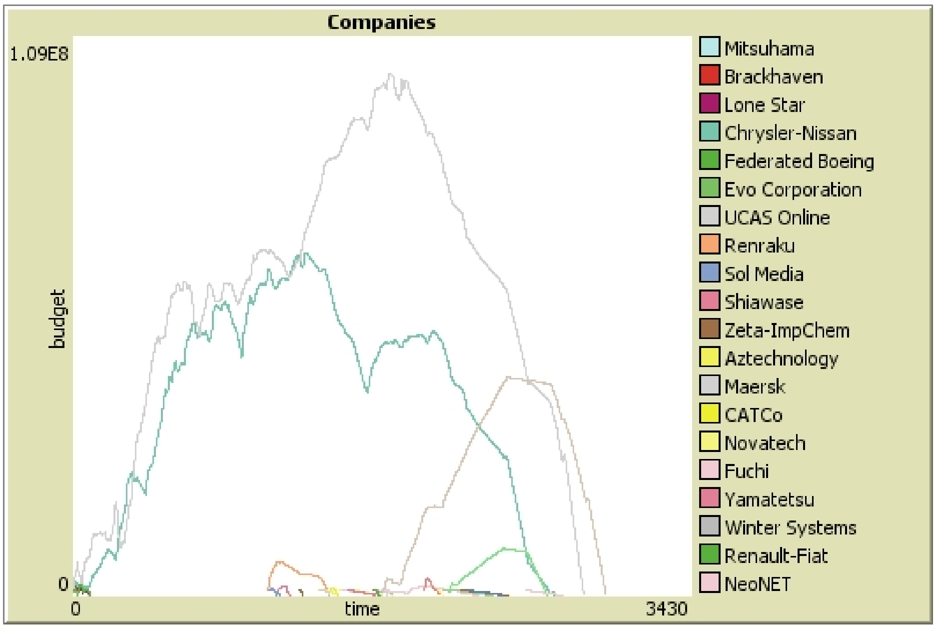

Figure 5.11: A typical run which contains most of the common behaviors that we can see in a lot of simulation runs.

Then, after some time, all companies are fading out and the simulation becomes lifeless. This does not necessarily happen, as there are also a lot of runs which apparently never end, but this is only because we cannot run them infinitely. As a matter of fact, the way the model is built it is almost bound to become instable at some point of time (or actually you could also call it stable after all companies vanished, as nothing will happen anymore). As the emergence and the decline of markets are both more or less completely random, there can always be a period were there is almost no market available, which obviously leads to decreasing company budgets.

Also in the long run with the deletion of unused processes there is at some point in time a period where there is such a market situation, that almost no company is researching for some while, and then knowledge is gradually dying out. When the process network is down to such a point, where no new combination is possible anymore, then there is no way left to avoid the breakdown. To show the amount of randomness which the model provides, we can see at figure 5.12 a completely different result with the exactly same parameters. However a lot of the elements which were discussed right now we can find here again, just in a slightly different shape.

Costs and revenues

Now it is about time to tweak some of the parameters. But there the first question that arises is obviously which parameters have the biggest influence on the simulation behavior? One group of parameters that is for sure very important, and which was already mentioned a couple of times

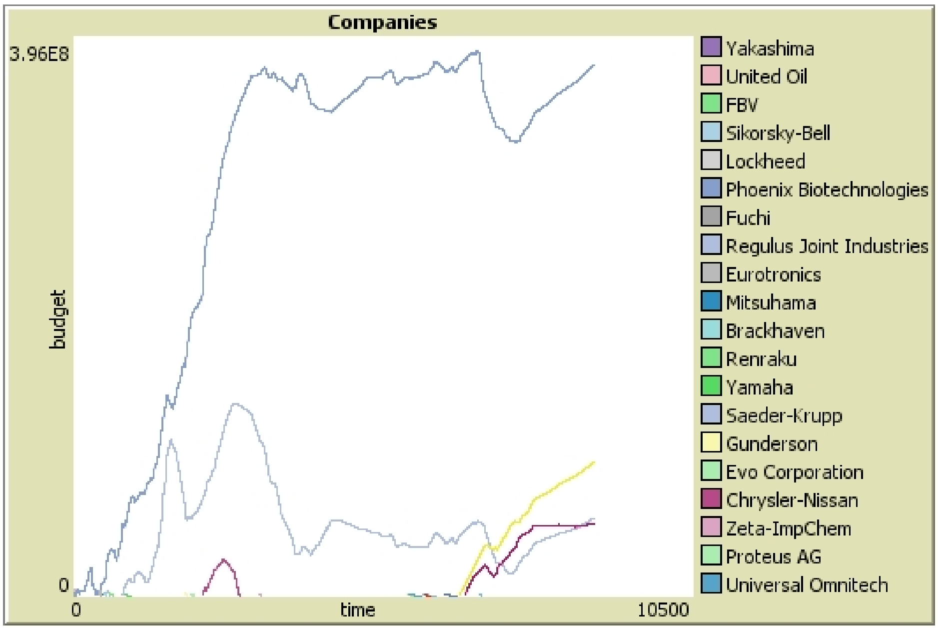

Figure 5.12: Another run with the exactly same parameters - while it looks totally different on the first glance, we can find again a lot of the described characteristics.

during the description of the parameters, is the whole cost/revenue relation. The min-fixed-costs are pretty straightforward: the higher you set them, the faster a company which is not performing is going down, but on a linear base. The same is true for the cost-per-research-percentage, with the exception that this can influence also top companies by limiting their growth.

In the 'best' case we see something like figure 5.13 where one company survives for some time, but with a lot of fluctuation and eventually the unavoidable extinction. In many other cases not even one company can survive the beginning phase. The other extreme with zero research costs is not so interesting in our implementation, as companies do not gradually adjust their research effort according to the costs - otherwise obviously it would have a bigger influence as more research would happen. But in our case it just leads to the fact that companies can rise even faster because of lesser costs, but it does not change the overall dynamics.

The most obvious parameters regarding costs and revenues are the avg-initial-costs and the avg-price. Usually a decrease of the price or an increase of the costs leads to a lesser number of profitable markets and a lesser average profit margin. Both means that companies will rise a little bit slower and are more dependent on the few profitable markets that exist, which means that they are stronger hurt in the case that one breaks away. But this is exactly the same as in the case that there are in general less markets available, which can be achieved with other parameters (which we will discuss later). If we choose for example average costs which are higher than the average price (let's say costs of 100/ price of 80) than there is only one thing that can be said about the result: It is highly unpredictable.

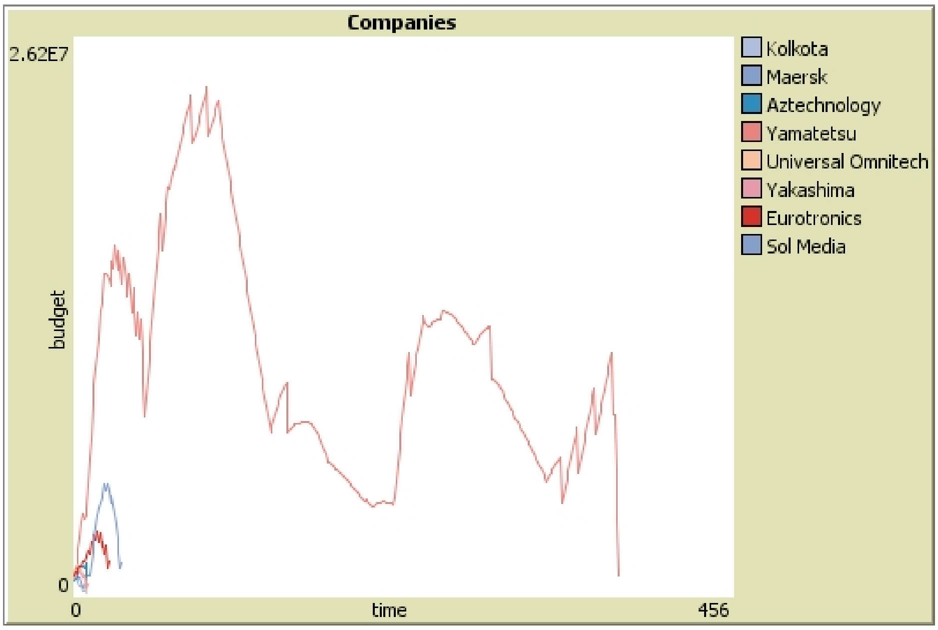

Figure 5.13: Maximum research costs (50.000 per percentage) lead to higher fluctuation and quicker extinction.

It seems that a constellation like that drastically reduces the predictability. Of course with all the randomness things can always look very differently, but with less (profitable) markets and a lower margin, the success of the companies is even more luck based. Often all die already in the beginning, while other runs look completely similar to our default setting. What is for sure is that the curves are steeper (the speed of decreasing budget is much faster) because it is more dependent on single markets. This leads to more chaotic looking results like in figure 5.14.

The last parameter pair in terms of costs that we want to analyze is years-production-fixed-costs and fixed-costs-factor-volume. The first sets the years of production which are taken into account for determining the 'size' of a company (and therefore its fixed costs) and the second parameter then defines the factor by which this production volume is transfered into costs. For example the default setting of 20/2.0 would mean the production volume from the last 20 years is taken and every produced unit leads to two cost units (Dollar, Euro, etc.). Obviously this method of defining fixed costs in a variable way has its disadvantages and also the actual values of these parameters are only relevant for this simulation and not in any real world context. However the relative level can simulate conditions from reality.

Let's take for example two extreme settings and compare them. At first we choose a combination of 100/0.5, which means that a change of size is only very slowly possible (we could think of expensive infrastructure, high amount of long-living but expensive machines or properties needed) but the actual costs per year are not that high. In that case we can see (like in figure

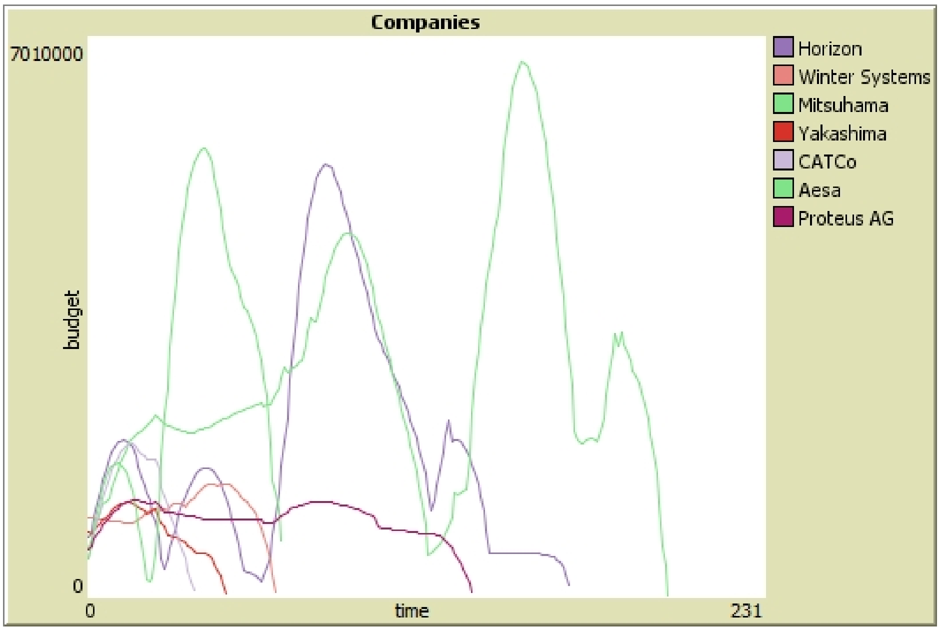

Figure 5.14: With the average costs being higher than the average price, more randomness and steeper curves are introduced.

5.15) that first of all, as costs are lower per year, in general more companies can survive.

Also, big companies might plummet quite fast in the beginning of a recession, but as their size shrinks, so does their fixed costs. Therefore we see a lot of absorbed ('cushioned') downfalls, which don't end with the extinction of the company, but with a rather fast turnaround. However, such a long duration of years which account for their fixed costs also means that there has to be some kind of invisible barrier at the top, which is almost impossible to break, as every increased production means more costs in the following years. As our markets can not grow there has to be some soft limit which can only be broken in times with an exceptional amount of markets available.

The opposite extreme setting would be a combination like 10/5.0. In this scenario in most runs companies don't survive too long. The downfalls are again very steep but in comparison to the previous setting, there seems to be no cushioning effect. Even with the size shrinking, the factor of 5 is apparently so high, that it does not change the trend in any significant way. What we also see (for example in figure 5.16) is that the peaks are very sharp, as the top values are unsustainable for any longer time. If we compare that to 5.15 we see that a high factor is definitely more influential regarding the extinction of companies than the amount of years, although obviously a combination of high values for both is even more deadly. A high factor seems to make it impossible for companies to regenerate from a slump, except they come up with an especially profitable innovation or a market is created which fits their knowledge.

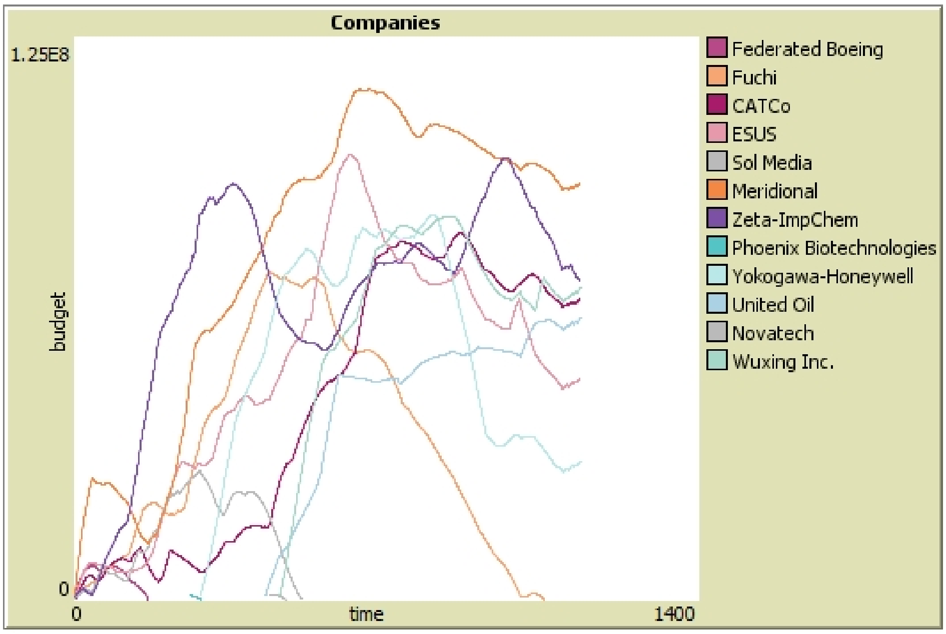

Figure 5.15: Maximum amount of years but a low factor for the determination of fixed costs leads to a lot of competitors and cushioned slumps.

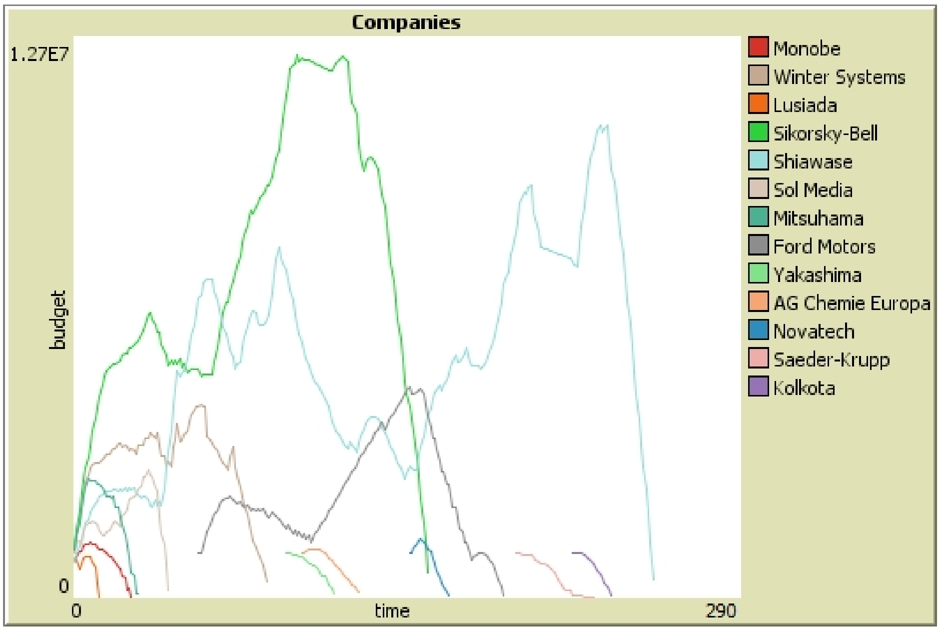

Figure 5.16: Low amount of years and maximum factor for the determination of fixed costs leads to quick extinctions and sharp peaks.

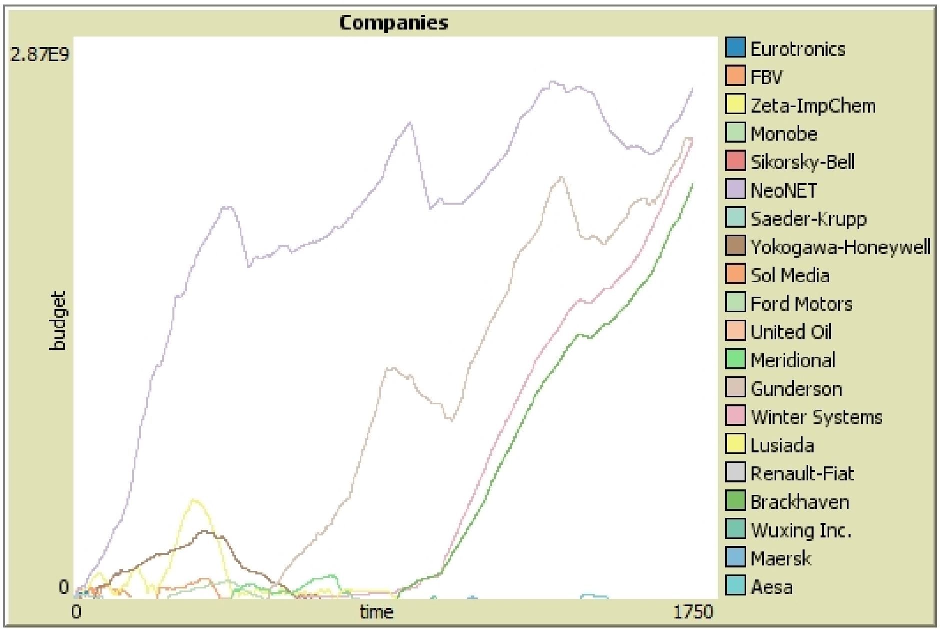

Figure 5.17: Maximum size of markets increases the peak level drastically and allows several companies to sustain this level.

Markets and competition

We treated the topic of costs so extensively because it is by far the most influential, and as we can see a big amount of parameters are related to it. However, to conclude this section, we will now leave the cost area and look at a few other parameters which also could have a big impact, even though we just can not cover all of them. Therefore the second topic we are looking into is markets, respectively their size, availability and competition for them.

The first obvious parameter to examine is avg-market-size. However, the results from changing this parameter are straightforward. If you choose the parameter too small (again it depends on other parameters what is considered to be too small), like for example 3.000 with everything else staying the same compared to the default setting, the simple result is, that all companies are struggling mightily and vanish very quickly. Even in the rare occasion where one survives the beginning phase, it goes extinct not much later. You would have to increase the number of markets, the possible profit and so on to give the companies even a chance to survive with such small markets.

For choosing the maximum market size (100.000) the results are not that trivial. Huge markets do not necessarily increase the success chance of an individual company - if they don't have the right products for the markets, they will still live rather shortly. What it does however is to lift the possible peaks to a much higher level, and as there is enough to share, several companies can be sustained at that level, as we can also see in figure 5.17.

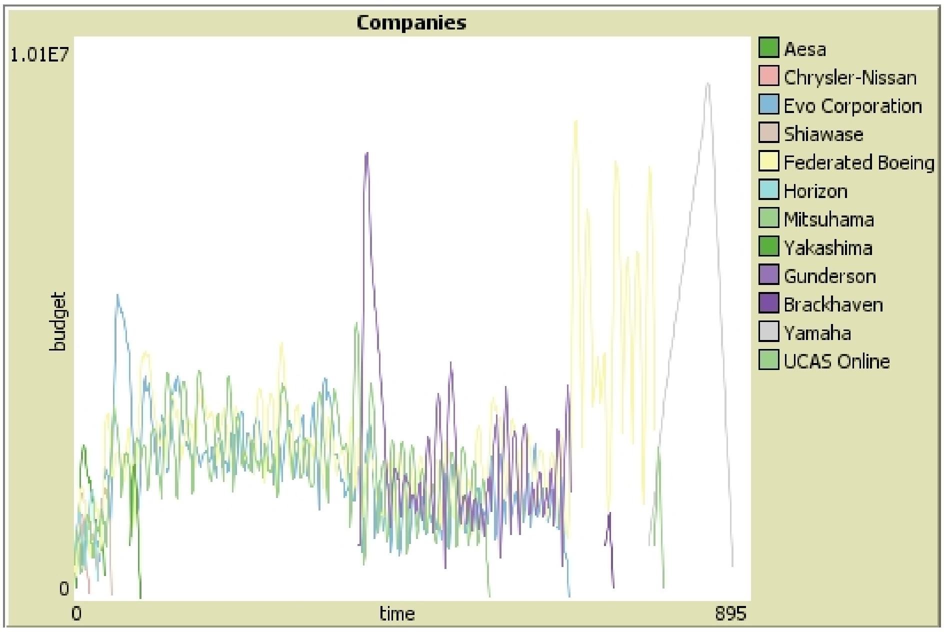

Figure 5.18: With a low acceptance threshold the budget becomes highly fluctuating.

Another relevant parameter is acceptance-threshold. This defines how 'open' markets are to different products. A high value means that only products which are very similar to the market taste will be sellable at all there. It has to be mentioned again that with the way the similarity is calculated, a value of 50 would already mean that almost every product would be accepted, therefore we have the default value of 85. If we use an actually quite low value like 70, the result couldn't be more different than what we are used to from the last figures. The company budget become highly fluctuating (see figure 5.18).

A likely explanation for this is that as almost every company can sell its products on nearly every market, everything is only driven by the taste changes of the markets, which determine the sold amount and price. Also if there are new competitors or if a company falls below the acceptance threshold (because of the taste change) and therefore out of the race, we automatically have a chain reaction which decreases or increases the share of everybody else.

The other extreme does not create such an interesting picture. With high values for the acceptance threshold (90+) either every company dies almost immediately because it can't serve any market or there is one which is lucky enough to fit a market in the beginning, which gives it a quick rise, but in the end it will also cease to exist very fast as soon as it falls below the threshold.

The next thing we should look at is the ratio between the number of initial companies and the number of initial markets. While an increase of markets obviously has a positive effect (and a decrease most probably a negative), it is not so clear with the number of companies. Less

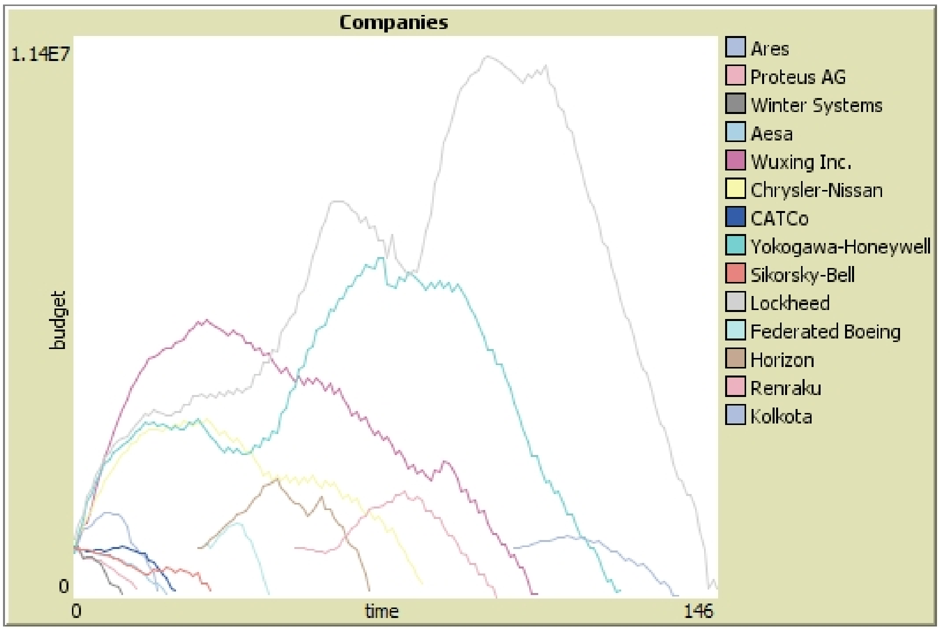

Figure 5.19: In the setting with ten initial companies and three markets many companies just don't grow enough to survive their first struggle.

initial companies means less competition, that's true, but there are also less potential candidates to survive, which leads to a picture maybe not necessarily expected: Namely, that the ability to serve a market is a bigger driver for survival than competition and therefore on average we will see more successful companies with a higher initial number than with a lower - as long as the number of markets stays the same and the markets are not extremely small. So now we will look what happens if we change the whole ratio. Again we look at the more extreme sides, as we know already the outcome for our default setting of seven companies and seven markets.

What happens if we change it to three companies and ten markets? Well, the effect of the increased markets is negligible, therefore there is also nothing really to show for in a figure. From the three initial companies on average only one succeeds (which fits to the ratio of two to three in the case of seven initial companies), if any at all. Often the run is already over before it even really started. If there are survivors of the beginning phase, and it comes to the events of creating new companies, the usual dynamics set in place and it doesn't look much different than in the default setting anymore.

The other way around it gets a little bit more interesting. The most noteworthy observation is that the low number of initial markets is relevant only for the very beginning. Because of so many companies who are all researching immediately a couple of new markets are created. Nevertheless we see that companies are struggling more in this setting than in the default scenario. Either they survive only barely and the usual mechanics then can set in place, or they go extinct even after they seemingly have survived the initial phase. Figure 5.19 is a good showcase for this finding. It seems that while eventually there are enough markets for several companies to survive, the high competition and the lack of markets in the beginning is still hurting them insofar as that even the successful companies cannot rise to a very high level. So when they then hit their first (almost inevitable) bump, they don't have enough budget to cushion that downfall, but are immediately fighting for their existence.

We will now conclude this section, as the space here simply doesn't allow to treat all parameters, but the readers are invited to experiment themselves of course. To give some ideas, we could for example not only look at the initial number of markets and companies, but also change their rates of creation (probability-new-company, similarity-lower-bound-new-market) and extinction (extreme-market-limit). Or we could change the amount of years until when a process becomes common knowledge (patent-duration) or irrelevant (duration-becoming-irrelevant). There is still a lot more to explore, but it should also be noted, that a lot of this exploration was already done implicitly while building the simulation and can be found in the recommendations for each parameter.

5.4 Economic Interpretation of the Results

This is now the time where usually questions are raised like: And what is the economic meaning of this simulation? What are the new insights that it delivers? But, as was stated several times throughout this work, it is not so much about the new insights that this model creates. I built a showcase model to demonstrate that the creation of novelty can be implemented in an agent-based simulation. This model is far away from picturing the economic truth, as we don't even have real market dynamics. The main insight should be, that this could be a starting point on how to model novelty in an evolutionary economic framework, and that it is feasible, but let's keep this topic for the conclusion.

However, we already saw that there are some patterns emerging in our runs, patterns that tell some kind of stories which seem to be remarkably similar to what we can observe in the economic reality. Often they occur for quite different reasons than in reality and some are probably even bound to happen because of the way the model is built. In this last section we will pick a couple of these patterns and try to interpret them.

Market structures

The first pattern or story that we want to look into is the market structure that usually arises in our simulation runs. To simplify things we are not looking into every single 'market' but rather see the whole simulation as one market and each 'modeled market' as one special product demand within this market - a view that is perfectly fine as we stated from the very beginning that our markets are just an abstraction to divide the demand, and this abstraction can be on any aggregation level. In our default setting (and many others) it seems that with seven initial companies, hardly more than two or three survive the initial struggle and that in general there are rarely less than two or three companies which are really on the top. There is no real perfect competition, but rather oligopolies are formed. Monopolies occur also sometimes in our simulation, but they are rarely sustainable as there are always new upstarts who challenge the reigning leaders.

This would fit pretty good to what we can observe in the real economy. It is well accepted in economic papers to treat international trade as an oligopolistic market form, or what we also find, 'monopolistic competition'. But, with that in mind, our eyes often tell us only what we want to see, and it could be easily possible that we just blank out all the runs which are not fitting to our expectation. Therefore I made an experiment with 100 runs of the default setting and the average number of companies which exist after 100 ticks and have more budget than the initial-budget parameter (and therefore are at least somewhat successful) are indeed 2,79 - which confirms the two to three we estimated.

However, as good as our model seems to fit to reality, it has to be clearly stated that this is not because it simulates the market dynamics so realistically - which it does not, as was already previously admitted. In our case it seems to be just the right mix of parameter values, especially the amount of knowledge each company has, the amount of markets that exist and the acceptance threshold that a process can be sold as product on a market. Previously we stated that when changing the initial-companies from seven to three, on average only one company does survive. While this is not completely true (our experiment with again 100 runs showed that it is on average 1,77) this indicates that these parameters result in some sort of ratio which is around 1/2 to 1/3 of surviving companies from the initial number, although as the number gets smaller, the ratio gets higher, as there will always be positive outliers with five or even more companies existing at measuring point of time.

Clusters of company formations

The second story we want to look into is the observation that new companies emerge usually in clusters, which means that there are long periods where no new competitors arise and then suddenly two or three within a short time. To undermine this impression, again an experiment was run to see the distribution of years (ticks) which are passed between two companies to emerge. If the observation is correct, we should see a spike around rather low values and the rest of the distribution should still contain some values, although we cannot expect a spike in the higher regions too, as the values would be too spread out.

Figure 5.20 seems to confirm our expectation. The spike in the low values is undeniably there, and is even stronger as it seems as the distribution classes are not equally big (the given numbers are the upper limits of each class). However regarding the spreading out to even higher values it has to be noted, that first of all the run limit was set to 1000 ticks. That means that for runs which would have gone longer, the possible values are not visible. Also this experiment was only measuring the distances between emergences that actually happened. Witch means that if until tick 1000 no new company emerged, nothing was measured, even if the distance would have been already quite big. If we consider these two shortcomings of our analysis we can say that actually there would be even more values in the very high region, and this is exactly what

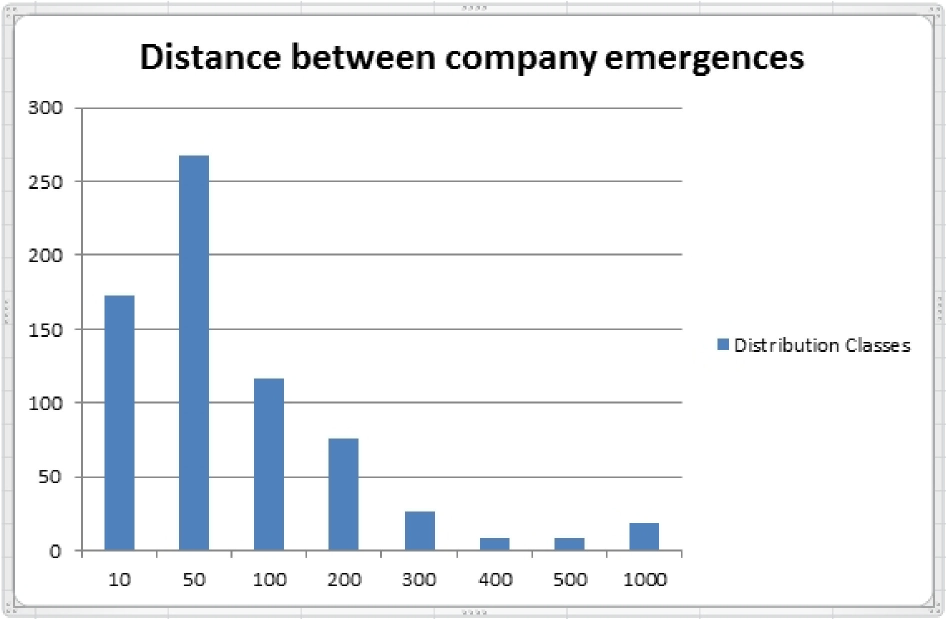

Figure 5.20: Distances between company emergences - There is undeniably a clustering taking place, considering the big amount of short distances.

we would expect, as for every wave of new companies, there are long times in between where nothing happens.

Now the question arises, is there any economic foundation for this phenomena? If we look at the reality, then our subjective impression would probably confirm this. When a new market opens because of technological innovations, we often see a mixture of some renowned companies and smaller startups which were especially founded for that market in competition. These new companies do not necessarily have to be in the technology race, often innovations just open the door for smaller businesses. One prime example is e-commerce: The establishment of the Internet as a business channel enabled a lot of people to do their business without the need of having a physical shop.

There is also enough of literature to back up this claim. In Shane [21] we find for example the following quote, referencing to previous research work:

> Research has shown that people are more likely to start new firms to exploit a new technology when the technical field is young (Utterback and Abernathy 1975) because the markets for new technologies are often too small to interest established firms, because established firms possess learning-curve advantages in mature technologies (Nelson 1995), and because the maturation of technology makes capital-intensity, the reduction of production costs, and scale economies important (Pavitt

and Wald 1971). ([21])

If we go back to our simulation, we find exactly this. As the criterion for new companies in our case is a drastically growing market, in most cases this will be a market which is young and was just established, otherwise these growth rates are hardly possible. The exception is a market which was not served by any company for some time because nobody had a fitting product - but even then one could argue that this is similar to a new market, even if the demand existed already much longer (and was already served once, long time ago). As this is an event which doesn't happen too often it seems a logical consequence that the emergence of new companies is clustered. But the same could be said about reality, and even though the reasons and the causal chain is much simpler in our model, in this regard it tells a quite good story.

The real drawback of our model however is that innovation has to come from established firms, as only afterwards new companies will emerge. This is only a part of what we see in reality, although some markets work like that: "Some industries have lower rates of new firm formation than others because established firms are the major innovators of new products and processes in those industries (Acs and Audretsch 1990)." ([21])

The rise and downfall of companies

The next 'story' is told rather fast, but important nevertheless. We want to look into patterns regarding the emergence and death of companies. We already said that competitors emerge mainly in bunches, but we didn't say anything about their survival yet. Obviously we just need to look at our simulation runs to see that most startups are not surviving very long. This is not surprising, as we know this fact very good from reality. However, there is no benefit in comparing the survival rate of our model with the numbers from reality. They will be completely different, for a whole bunch of reasons. For example a real economy consists of much more competitors than in our case, therefore there will be much more startups and consequently a higher failure rate.

In our model we need at least two or three companies surviving to make an interesting run but at the same time we cannot afford to start with 100 companies as it would just unnecessarily slow down the simulation. There is also no comparability regarding the time frame. We randomly declared that one tick in the simulation is one year, but this is an arbitrary choice, which just makes the model easier to comprehend. But there is no reason that it might not be more close to real figures if we define a tick to be a day or a month - it is just one time unit and therefore not comparable to reality. After all it is just a model, and therefore some abstraction is necessary.

However there is one thing we can compare, and that is the survival rate between those companies which are there from the beginning (which we already measured) and startups which join later. The expectation (at least for me) would be that startups which join in the middle of the race have a lower survivability, as they usually join a crowded market where there are already some leaders. On the other hand they have the advantage of being automatically able to serve at least one upcoming market, so maybe this is equaling things out or even giving them better chances.

The result is somewhat surprising: Among our initial companies we had an average survival rate of 0,399 (with default settings) and now we have a survival rate for startups of 0,426. While this is only slightly higher (and might be statistically insignificant), it for sure is not an disadvantage to join the race later in our model. It looks like the advantage of having for sure one market to serve is outweighing the drawback of facing an established competition (which additionally might not even exist in some runs, as there is a good chance that most competitors already vanished).

Another observation that was already analyzed while describing the general patterns of the simulation runs, is that there is no run that is lasting forever. Obviously this is mainly due to factors which are model-immanent, but the notion that no company is surviving forever seems to be important. Of course our data from reality is very limited and the big international corporations that we know nowadays barely exist longer than half a century - at least in the size that we know them nowadays. However it would be interesting to theorize if there is some limit of how long it is realistic to expect a company to survive in a somewhat stable environment (no wars, no radical change of the foundations of economics like the introduction of capitalism and so on). Personally I would expect that every company vanishes at some point in time - there are just too many 'opportunities' where a company can fail, like general economic crises, the misjudgment of new technologies, the inability to keep up and so on. And in this regard I think that this model is doing a good service at pointing out how many risks there are for a company and showing the different ways of failure, so to say.

Social learning and common knowledge

It is not really a story or a pattern, but rather an observation that in order to create interesting dynamics, we have to set the duration until a process becomes common knowledge rather high. With low values very fast all companies look more or less the same, in their knowledge network as well as in regard which markets they are serving - as a result also their development is pretty identical. Therefore in the default setting the whole process that processes become common knowledge is more or less deactivated. In general, this result is not particularly surprising, as Chang and Harrington [3] came to similar results:

> When the communications technology is poor, it was found, not surprisingly, that technological improvements enhance performance. What was surprising is that if the communications technology is sufficiently effective, further improvements are detrimental. Better communications allows more social learning among agents and this endogenously results in a more structured network as each agent is able to identify those agents from which she can learn. (As agents faced different environments, an agent wants to connect with those agents who face similar environments and thus are likely to have applicable solutions.) This, however, has the unfortunate

> by-product that it results in agents having very similar solutions and the ensuing lack of diversity within the social network meant that the population of agents is ill-equipped to adapt to a changing environment. Thus, a better communications technology can lead to too structured a network from the perspective of promoting innovation. ([3])

While there is no real communication between our agents, the principles are the same. In our case it is even more extreme as every company learns from all others as soon as the process becomes common knowledge. Therefore a more advanced approach would be desirable to cover this whole topic. The authors present in the same paper a model where the agents have the choice to either pursue individual learning (research) or social learning (observing others). This requires some sort of communication, as the social learners have to know which competitors they should observe.

Without a doubt this would be a very interesting extension for our model. Especially as the key finding of Chang and Harrington [3] is, that companies specialize over time to become one of both learning types. In our case this could provide a possibility to bring back the idea of processes becoming common knowledge, but in a more strategic, intended way, controlled by the individual companies. As a positive side effect it would also further heterogenize our agents, which anyway would be a big goal for further expansions.

The (dis)advantage of researching

The last question we want to address is if companies which are researching more than average are having a competitive advantage or not? We already have to realize in advance, that this is a tough question to answer. What we can say for sure is that researching is not a guarantee for success. On the other hand a company that will never research will automatically drop out of the race as soon as the markets which it serves are vanishing (at least in our default setting where established processes barely become common knowledge). Analyzing if there is any correlation between the amount of research and success has to be done with a grain of salt as in our implementation struggling companies will usually research more, as the conditions for successful companies for research seem to be stricter than for struggling ones. So most likely we reverted the causal chain here. However, it is still interesting to have a closer look at this, as maybe our assumption is even totally wrong and there is no correlation between the research percentage and success in terms of survival at all.

In the next experiment I measured the research percentage and the corresponding survival duration. We can see the results at figure 5.21. There are a few comments and explanations necessary to put this picture into place, however I hope to not over-interpret the results to fit my opinion. Obviously the reader anyway should make his/her own mind. The first we can see is that there is a convex limitation curve which makes certain combinations impossible (except some outliers which were probably calculation errors). This is because during the first 50 years, every company is researching automatically because of the public funding. That means that a company

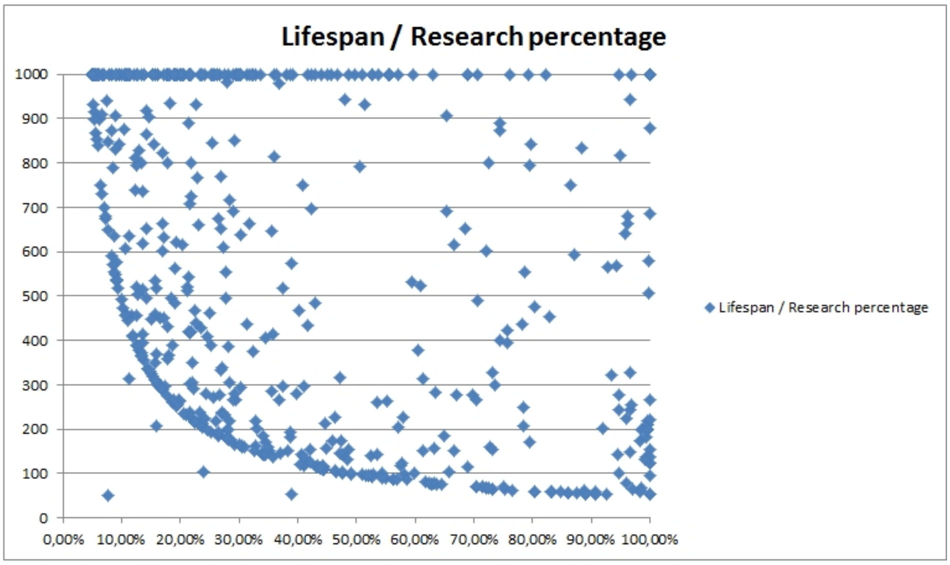

Figure 5.21: Correlation between lifespan of a company and research percentage during that lifespan.

with lifespan 50 has automatically a percentage of 100% while with a lifespan of 1000 there has to be at least a percentage of five percent.

Other remarks have to be made: The maximum run length was 1000 in our experiment, so this is the reason why we find a lot of data points there, these are the companies which existed from the beginning and lasted to the end. We obviously don't know how their performance would have continued, but the research percentage distribution between them is of course still telling. On the other hand we see a couple of companies with a lifespan of over 200 years and still a research percentage of 100%.

While it is possible that a company is first researching because it is so successful and then immediately afterwards continues researching because it can't serve any market anymore, this combination is rather unlikely. More probable is the scenario that these were very successful startup companies which were rising all the time and then the simulation run was just over because of the limit of 1000 years. Again we don't know how that company would have performed afterwards, but for sure it would have pushed its data point more to the upper right corner with every additional year it would have survived.

So, while that rather empty space on the upper right corner (high research rate with high durability) would be a little bit less empty, we can say almost with certainty that there is no positive correlation between the research percentage and the lifespan. If we find any correlation, than it would be rather a negative one (the correlation coefficient is -0.39).

However we also shouldn't take this insight too serious. My personal interpretation is that the criteria for enabling research are rather strict in our model, and as every company has a lot of ups and downs during the simulation runs, it is almost inevitable that the longer the company exists, the more stretches exist where it doesn't meet these criteria. Therefore it is almost impossible to sustain a research percentage over 50%, although some companies are still able to do it. Also it is not so easy to be successful for 10 or more consecutive years, so easing the condition of parameter years-of-growth-for-research might also increase the average research rates of successful companies.

Honestly, I don't see any implications or relation for/to the real economy here. One should not forget that this result could look already completely different with simply decreasing the research costs. Also most probably if we would do the same experiment, but compare the average budget of a company with the research percentage, this experiment would result more in favor of researching. But this is only speculation on my side and proving it would go beyond the scope of this work. Anyway it probably wouldn't give much additional insight, as the causal chain here again would be that we are defining if (un-) successful companies are researching and an analysis would not necessarily tell anything the other way round. If there is any statement that the model is telling us, then it is the triviality that we knew already before: Researching companies are not necessarily more successful. However I still strongly belief that in reality there is a positive correlation, so the main conclusion that we can draw from all this is that there is still work to do to improve the model in that regard.

CHAPTER 6

Conclusion

In the course of this work we looked at a lot of different ideas coming from various theoretical streams. Therefore I want to recapitulate the most essential thoughts from this journey through scientific fields and time periods. Only with recalling our starting points we can draw conclusions about the actual outcome of this thesis.

We started with the recognition of technological progress as one of the main motors for economic growth. We therefore looked into traditional growth models, were we stumble upon the famous Solow residual and the rather startling fact that technological progress was treated as an exogenous process in the Solow model. The other two classical models we encountered (Arrows and Shell) were at least endogenizing the technical progress, but still neither of them was explaining what knowledge really is (it is mainly treated as another quantifiable resource) and how it is formed. It is just assumed to be the product of some process, for example of the investment activities of the private sector. Nowhere is explained how this process works and what innovation means for companies. Only the Shell model at least depicts the trade-off between technological progress and manufacturing, by having two distinct sectors which share the same input resources.

We then took a step away from the neoclassical view to a more evolutionary spirit, which is Schumpeter and his entrepreneur. Here we find a lot of good stories, but no consistent theory. He criticizes the lack of room for development within those neoclassical models, the room for evolution in the world view of general equilibria. In his view for small changes, where only the exogenous data is changing (flavor of the consumers, events outside of economics), the traditional models are sufficient. “But, these means are failing (...) where the economic life itself changes its own data suddenly.” ([19], translated by the author) So, for the big technological revolutions, in his opinion the neoclassical models have no answers. Again, I want to highlight the two most outstanding concepts of Schumpeter regarding innovation:

- Innovation is **not driven by the consumer**, but happens in the sphere of the industrial and commercial life. This means that **novelty is created by firms**.
- Innovation is created by the usage of **new combinations** of production means.

While his ideas around the entrepreneur might be considered outdated, his views regarding business cycles are not. The rise and fall of companies is according to him highly related to innovations. In his radical idea, he wants that we suppose that every innovation is embodied in a new firm founded for exactly this purpose. Interestingly enough, in the age of the Internet this might be more true than it had been over the centuries before. Even if we argue that a company is able to produce more than one innovation, this does not change the view that companies which are not able to present new products are supposed to lose the race and therefore collapse. This picture couldn't fit more perfectly to the views of evolutionary economics and of course also to our model, as we now already know.

We then went back to neoclassical economics and looked intensively at a modern approach by Aghion and Howitt. This model incorporates interesting ideas and is even trying to implement Schumpeterian ideas like the business cycles. Still the basic points of critique remain the same, as a neoclassical, mathematical model is not able to represent real change and dynamics, even less able to explain the process which is making this change emerge.

Not very surprisingly we therefore turned to evolutionary economics. During the attempt of defining what evolutionary economics is, we find the very interesting proposition from Potts [18] that in our economic model the set of connections between the elements should have the same importance as the elements themselves. As we already know, this is very crucial for the model of this work, as connections according to this definition have their own properties and the combination and recombination of elements is strongly dependent on the connections of these elements.

This proposition leads to another consequence: The field concept which is used in mathematical economic models cannot deal with this concept of connections and their non-existence between certain elements. Only graph theory and complex systems can. This coincides with the assumptions of this work, namely that a mathematical model is not able to represent the creation of novelty in the way that we imagine, and that this is only possible with a complex system, like our agent-based model is. It is also not by accident that knowledge (the processes and their connections) are modeled as a network in our implementation, but very much inspired by Potts.

In our sidestep to novelty in biology we recognize that we have to be careful with drawing too many analogies, but that there are nevertheless parallels. In both fields we have both, calm (regimes) and very revolutionary phases. In my opinion there still remains the difference that in biology even within these revolutionary times every step itself is as small as possible and 'random' while this is not true for innovations, as there the process of creating novelty involves the human mind and therefore entire (re-) combinations of old elements.

In our search for novelty in evolutionary economics, this work was also heavily influenced by Witt. Putting his critique of the neoclassical approach aside, which he shares with other authors, he states that novelty is by it's very nature unpredictable and connected to uncertainty - A fact that makes it hard to deal with in any theoretical model. The main problem that traditional economics has with this fact is that uncertainty and risks don't really have a place or would be avoided according to these theories. However, speculative investment in R&D is totally rational in the real world and therefore should also be in any economic model - and it is in evolutionary economics.

The last thing we did before actually building our model was to look in existing literature how novelty can theoretically emerge and how this could be modeled. Witt [24, 25] was a great inspiration for this. Just randomly combining old elements will not lead us to any innovation. There has to exist familiarity or at least some kind of connection with the context to which the element is brought, so that this element, which is coming from a different context, still has some meaning. "Only if the conceptual inputs to be blended have some kind of similarity - some common element - can their non-common elements be integrated in an associative act into a new, meaningful concept." ([24])

Witt defines a generative and an interpretative operation and states that the generative operation "can, in principle, be done by any arbitrary device." Finding sense in these new combinations or aligning the multiple contexts that have been mixed, however, is something completely else and can, from a current point of view, only be done by a human agent. However, the human mind usually does both operations at the same time. In our model the generative operation is indeed automated. To kind of skip the interpretative step we use the 'trick' to abstract the generated concepts from any real meaning and just give them quantitative properties that could mean anything.

For the process of (re-) combination Witt introduces the concept of *propositional networks*, which is used to describe the sharing of a common concept. For simple concepts like sentences this could for example be sharing the same subject. While the generative operation can be carried out by a machine, to avoid nonsensical statements, it is not recommended to reshuffle elements just randomly. What is effective on the other hand is to use the commonalities, the similar elements, as a starting point for crossing over the non-common elements.

Therefore, the generative operation is far from a "blind" variation, but there is a strong guidance towards what is chosen to be recombined. In the current implementation of our model however this guidance does only exist in the form that companies prefer to use concepts which they already used for combination successfully before. It would be one of the most important starting points for an extension to make them actually look into the properties of the concepts and choose their combination based on their similarities - in the best case even already with a certain outcome in mind.

Witt puts up several hypothesis, whereof some of them were already mentioned indirectly or summarized in this conclusion. But especially number 4 and 5 should be repeated here in hindsight, to see that the model which resulted from this work was especially trying to incorporate his theories:

> Novelty is revealed through inductive operations whose outcome, by definition, cannot be derived from the premises. Carrying out the inductive operations requires time, and thus prevents the meaning of novelty being instantaneously accessible.
>
> By re-using newly created concepts in further iterations of the generative operation, an infinitely growing number of concepts can emerge from a finite number of initial elements, provided that the share of combinations to which new meaning can be attributed is non-vanishing. ([25])

Obviously this work drew inspiration also from other papers, of which Hanappi and Hanappi-Egger [14] should be named. Especially the picture of having a generator and some kind of survival gate as a testing mechanism to determine which elements are successful proved to be very useful. We can imagine the ‘invasion’ as the introduction of new products and the survival gate as our market feedback, whereas these gates are changing all the time, as they can either ‘move’ (change of taste) or complete disappear and emerge (creation and deletion of markets).

For the creation of new elements they came up with the idea of **genetic algorithms** as a solution: “In a genetic algorithm the behavior of a type is represented as a bit string, and one could think of the innate growth rate of this type as its fitness function. The crucial assumption then is that parts of the bit strings of successful types can be glued together - and the new fitness of this newly born type is immediately available.” ([14]) In our model we use property vectors instead of bit strings, but the idea is very similar.

After we see now how much theoretical work and ideas have influenced the model of this master thesis, it shouldn’t be very surprising that there is still a lot of room for possible enhancements. The groundbreaking insights as a result of this model which everyone would always like to see are limited. The most revolutionary aspect of this work is that it actually tried to put all these ideas together and for once build a coherent model that can even be run as a simulation.

In this regard this work can be seen as a real success. It used the concept of (re-) combination as an generating engine and implemented the elements used for it (the processes) in the form of a knowledge network, where not only the elements themselves, but also the connections are relevant. A possible extension would be to give those connections even more meaning, by introducing a more strategic behavior of the companies and giving the connections some influence on the predicted result of a combination instead of just defining if a combination is possible at all or not, like it is the case for now.

Also in the area of how knowledge develops there are a lot of ideas for improvement. We showed the concept of iterative upgrades, which (re-)combine already combined elements in a more effective way as a possibility to have real technological progress over time in our model. Another idea in that regard is to introduce supportive processes, which don't change the outcome of the dominant process significantly, but make them more productive. Coming from a different side (Chang and Harrington [3]) we could improve the learning process by introducing mechanisms to develop specialists in either individual or social learning, which could also reduce the problem that too much common knowledge creates to homogeneous agents.

We do have a testing mechanism with the markets that defines which elements are successful and which not. However obviously with the goal to be used in a real economics model, one of the aims has to be to introduce real market dynamics. Nevertheless the main aim of this model was to be the starting point for others; in best case even a template, but at least a source of motivation to leave the pure theorizing behind and start to build actual models. Only by providing concrete solutions the school of evolutionary economics can become a real alternative to neoclassical theories. Well told stories alone will not be enough.

Nevertheless, as crude as the model here might be in many aspects, it should not be underestimated in terms of the dynamics that it can provide. By making research and the emergence and death of companies and markets endogenous processes of the model, no run is ever like any other. We also showed that many dynamics that we see in the real economy are also happening in our simulation. Some of these phenomenas are intentional and just a result of the model restrictions, while others emerged rather surprisingly.

I firmly believe that with exploring the behavior space even better, already this basic version of the model could resemble the reality even closer. One can only imagine which wonders all the mentioned possible extensions could do. A full market implementation together with more strategic options regarding the research decisions could raise this model to a whole other level of expressiveness. Therefore this is a call to give it a try, to either extend this model or to come up with complete new implementations, so that we have a whole pool of possible solutions on how to model the emergence of novelty. A concept which is not only vital for the rise of evolutionary economic theories, but for the struggling discipline of economics as a whole.

Bibliography

[1] Philippe Aghion and Peter Howitt. *Endogenous Growth Theory*. MIT Press, Cambridge, MA, USA and London, England, 1998.

[2] Brian W. Arthur. Out-of-equilibrium economics and agent-based modelling. In Kenneth L. Judd and Tesfatsion Leigh, editors, *Handbook of Computational Economics, Vol. 2: Agent-Based Computational Economics*, pages 1551 – 1563. North-Holland, 2006.

[3] Myong-Hun Chang and Joseph E. Harrington. Individual learning and social learning: Endogenous division of cognitive labor in a population of co-evolving problem-solvers. *Administrative Sciences*, 3(3):53–75, 2013.

[4] Charles Darwin. *On the Origin of Species - By Means of Natural Selection*. J. Murray, London, England, 1859 (2006).

[5] Richard Dawkins. *The Blind Watchmaker*. Penguin Books, London, England, 14 edition, 1986 (2006).

[6] Kurt Dopfer. The evolutionary paradigm of economics. Originally supposed to be the Introduction of the collection 'Evolutionary Foundations of Economics', 2003.

[7] Kurt Dopfer, John Foster, and Jason Potts. Micro-meso-macro. *Journal of Evolutionary Economics*, 2004.

[8] John Foster and J. Stan Metcalfe. Economic emergence: An evolutionary economic perspective. *Journal of Economic Behavior & Organization 82*, 2012.

[9] Gene M. Grossman and Elhanan Helpman. *Innovation and Growth in the Global Economy*. MIT Press, Cambridge, MA, USA and London, England, 7 edition, 1991 (2001).

[10] Hardy Hanappi. Evolutionary economic programs. Version 12-11-2003, paper presented at the annual meeting of the Ausschuss für Evolutionäre Ökonomik, Erfurt, July 3-5, 2003, 2003.

[11] Hardy Hanappi. Schumpeter. In Michael Dietrich and Jackie Krafft, editors, *Handbook on the Economics and Theory of the Firm*. Edward Elgar Publishing, 2012.

[12] Hardy Hanappi. Evolutionary political economy in crisis mode. Version 03-09-2013, paper contributed to the special issue 'Frontiers in Evolutionary Economics' published in the Journal of Economics and Statistics (forthcoming), 2013.

[13] Hardy Hanappi. Evolutionary dynamics in revolutionary times. In Sebastiano Fadda and Pasquale Tridico, editors, *Institutions and Economic Development after the Finanical Crisis*. Routledge, 2013.

[14] Hardy Hanappi and Edeltraud Hanappi-Egger. New combinations: Taking schumpeter's concept serious. *Munich Personal RePEc Archive*, 2004.

[15] Geoffrey M. Hodgson. Darwinism in economics: from analogy to ontology. *Journal of Evolutionary Economics*, 2002.

[16] Stuart A. Kauffman. *The Origins of Order*. Oxford University Press, Oxford, England and New York, USA, 1993.

[17] Stuart A. Kauffman. *At Home in the Universe*. Oxford University Press, Oxford, England and New York, USA, 1995.

[18] Jason Potts. *The New Evolutionary Microeconomics*. Edward Elgar Publishing, Cheltenham, England and Northampton, MA, USA, 1 edition, 2000.

[19] Joseph Schumpeter. *Theorie der wirtschaftlichen Entwicklung*. Duncker und Humblot, Berlin, Germany, 6 edition, 1911 (1964).

[20] Joseph Schumpeter. *Business Cycles (Konjunkturzyklen)*. McGraw-Hill Book Company Inc., New York, USA and London, England, 1939 (1961).

[21] Scott Shane. Technological opportunities and new firm creation. *Management Science, Vol. 47, No. 2 (Feb., 2001)*, 2001.

[22] Herbert Simon. *The Science of the Artificial*. MIT Press, Cambridge, MA, USA, 1969(1985).

[23] Ulrich Witt. Evolutionary economics and the extension of evolution to the economy. Paper to be presented at the Workshop on Evolutionary Economics, Schloss Wartensee, Switzerland, 2003.

[24] Ulrich Witt. Novelty and the bounds of unknowledge in economics. *Papers on Economics and Evolution, No. 0707*, 2007.

[25] Ulrich Witt. Propositions about novelty. *Journal of Economic Behaviour & Organization 70*, 2008.