Agent-based Computational Economics: Exploring the Evolution of Trade Networks

Gemma Hagen
⟨gh206@doc.ic.ac.uk⟩
Department of Computing
Imperial College London

Supervisor: Abbas Edalat ⟨ae@doc.ic.ac.uk⟩
Second Marker: Francesca Toni ⟨ft@doc.ic.ac.uk⟩

June 26, 2009

Abstract

As the field of economics seeks to further its understanding of the links between the micro and macro sides of the area, economists are becoming increasingly appreciative of agent-based models. Cross-silos approaches to problems are providing for a new way of thinking about how to explain complex phenomenon in science. Agent-based modelling provides a novel and effective way of explaining how such complex systems arrive at macro phenomenon through attempting to grow said societies. The project that I am undertaking is an extension of the work completed by Allen Wilhite which is documented in his paper "Bilateral Trade and Small-world Networks." His aim was to explore the efficiency of various trade network formations through an agent based simulation. In his work agents are able to produce or trade one of two goods. The report documents the extensions to this model in order to explore the evolution of trade networks and traits of agent behaviour. Furthermore, dependencies between agents are examined as well as wealth distributions. The extensions are evaluated both in the context of the simulation and the real world.

Acknowledgements

It is with great pleasure that I acknowledge the help, guidance and support of my supervisor, Prof. Abbas Edalat. Without his critical assessment of my methods, and willingness to provide both support and suggestions for solutions to problems faced, I would not have been able to achieve the results displayed here. I especially appreciate the brainstorming sessions which furthered my insight into the complexities of human interaction, and significantly shaped the direction of my project.

I would also like to thank Dr. Francesca Toni, my second marker for reviewing the progress of my project and giving me pointers to current research and platforms for implementation of agent-based simulations. In addition, her challenging questions permitted me to develop a greater understanding of the model at hand.

Contents

Chapter 1

Introduction

This report documents an investigation into the applicability and credibility of agent-based modelling to the study of economics. A simulation engine is created, using a fast and fairly rigid Java core, supplemented by an easily extensible scripting layer, with a simple web interface allowing customisation of operational parameters and easy deployment of simulations across multiple machines. Basing the core model on the work by Allen Wilhite, a world in which agents are able to produce and trade one of two goods is developed. Various network topologies restrict with whom agents can trade and, in addition, through whom agents can search in order to find a trade partner. Extensions are added to give agents the opportunity to learn about suitable trade partners through remembering exchanges, and a further type of exchange is permitted - the exchange of the knowledge of the existence of agents in order to grow a global economy from the bottom up. This can be considered to be analogous to a simplified model of human communication in business. In addition, investigations into the effect of necessary consumption are added to enhance the realism of the model and permit the analysis of bankruptcy of agents. This enforces agents to consume goods, encapsulating the finite nature of the resources, for instance fuel, upon which all businesses in the real world depend. Finally, a genetic algorithm is employed in order to gain insight into what makes a wealthy agent, which highlights both the successes and limitations of the model adopted.

As extensions are added, the effect on the dynamics and the evolution of the simulation is thoroughly evaluated, both in the context of how the model is improved, and with respect to how well findings fit with economic theory.

Economics is currently receiving criticism, especially given the current recession, from both academics outside the field and economists themselves. As the world becomes more globalised more interdependencies arise, leading to complexity that the models and assumptions of Neoclassical economics are struggling to address. Economics as it is today is becoming increasingly outdated. There exists a need for change. As cross-silos approaches to complex problems gain pace in the academic world, and research into complex systems is furthered, it becomes increasingly apparent that the economy is not as simple as it once was. It is necessary to revitalise economics, bringing it up to speed through adopting perspectives from complex systems, and studying it as such. This entails embracing the notion of emergence, which describes the process by which macro phenomena emerge from micro interactions. In addition, focus lies not on the elements composing the system, but rather on the interactions between these elements. Agent-based modelling seems like a powerful tool to facilitate developing models and growing societies, allowing examination of how micro interactions can result in these macro phenomena.

This report is illustrative of the fact that agent-based modelling has the potential to provide huge insight and compelling results, facilitating the explanation of the dynamics of the billions of people interacting to form the economy. In addition, by requiring that agents do nothing but strive to further their position in the economy by maximising their utility, many complex phenomena in economics are witnessed. Agents begin to specialise in whether they produce or trade, and traders become loyal to their partners. Wealth distributions are analysed and, as the model is developed, become increasingly realistic. That the rich get richer and the poor get poorer is shown not to be solely a saying, as wealth condenses on to the already wealthy. Learning facilitates the trade of knowledge of agents, and in a world of free trade without barriers, globalisation emerges and demonstrates its efficiency as prices converge, and creative destruction is witnessed as the most effective agents rise to become bridges across districts. In addition, the architecture of this globalisation is examined and compared to that of the real world.

1.1 Key Contributions

- The implementation of an extensible economic simulation, leveraging modern scripting language paradigms which facilitate rapid and simple development of complex logic describing the way in which agents interact. *(Section 9.2)*
- A simple web interface, enabling fast addition of configuration options, permitting easily customisable simulations.*(See Section 9.4)*
- The evaluation of results attained from the simulation, in both an economic context and in the context of the performance of the model. *(See Sections 3.3, 4.2, 5.3, 6.3, 7.3, 8.3)*
- A critical analysis of potential limitations of the model adopted. *(See Section 10.1)*
- Through results obtained, suggestions of the importance of network structure in the economy, as well as suggestions for direct applicability of these ideas are made. *(See Chapters 4, 5)*

The report consists of a description of Wilhite's model and a description of the motivation for and implementation of extensions, followed by a thorough evaluation for each. Table 1.1 summarises the motivation for extensions.

Extension	Motivation
Restriction of the percentage of agents who can produce both goods.	Wilhite's model inadvertently prohibited large amounts of trade, and a larger volume of trade was necessary to soundly evaluate the formation of trade networks.
Forcing agents to consume one of the goods.	Intended to capture the fact that business in the real world incur a need for resources to facilitate production (for instance factories) and trade (for instance shipping costs), and to analyse dependency in networks through investigating "bankruptcy chains".
Permitting agents to remember exchange.	Intended to investigate the usefulness of memory to agents, specifically traders, and to better reflect the common occurrence in reality of returning to the same source to buy or sell through having learned that it is least costly / most profitable.
Adding exchange of the knowledge of the existence of other agents.	Intended to attempt to grow a global economy, and to overcome the limitation in Wilhite's model that agents acting as "bridges" for trade between districts were chosen at random.
Addition of a genetic algorithm to evolve agents	Intended to gain insight into what makes a wealthy agent, given the initial conditions of the simulation.

Table 1.1: Summary of extension with corresponding motivation

Chapter 2

Background

2.1 Economics

Neoclassical economics is the dominating force of economics today. It steers the decisions of governments in making economic policies, and of companies in mergers and acquisitions. It commands text books of A-Level and undergraduate studies and its influence has been hugely beneficial for over a hundred years. All of the G8 countries have avidly followed the theories and solutions it has introduced, and are obviously extremely wealthy from doing so. However, over the past decade, there has been increasing criticism over the robustness of these theories. People's concern over the correlation of these theories to the real world is growing and companies with high stature, such as General Electric, have even closed down their economic departments. This is not to say that the theories are useless, moreover, there exists a progressive consensus that the field is not fulfilling its potential as a science - that economics is in need of a makeover. This is not just a wave of concern coming from scientists outside the field, but economists themselves are engaging in self critical analysis of the area. Joseph Stiglitz was quoted in an article in *The New Yorker* by Cassidy [5] "Anybody looking at these models would say they can't provide an accurate description of the real world."

2.1.1 Economic History: An Extremely Brief Overview

Obviously an account of economic history is hardly feasible, however, it is interesting to see the development of the field - or at least the key turning points. A good place to start would be with the revelations of Adam Smith. Adam Smith was a moral philosopher turned economist who brought huge ideas to, at the time, the underdeveloped field. Smith attempted to answer two important questions, namely how wealth is created and how it is allocated. For the first question, he argued that labour productivity, which was largely dependent on the division and specialisation of labour, is what creates wealth. It is the act of using raw materials to create items that other people want in an efficient manner. For the latter, he claimed that if people were free to trade on their own, the pursuit of an individual's self interest would drive them to provide the resources people want at the price they are willing to pay.[17] This came to be known as the *invisible hand* and is a good example of emergent behaviour since through an individual's pursuit of self interest they would be inadvertently benefiting society as a whole, a macro phenomenon, through efficient resource allocation. Smith explained that the key mechanism in these competitive markets was the price at which people trade. This lead to the idea that the market clears at the point where supply is equal to demand - a core concept in neoclassical economics today.

Although this was a profound idea, economics still massively lacked mathematical rigour. It was Walras, along with his compatriots, who strived to find a way of bringing this element to the table. His conjecture was that the equilibria witnessed so frequently in nature was analogous to that of markets clearing.[27] The notion of equilibrium however is more complex than simply opposing forces coming in to balance as there are many sorts of equilibrium apparent in different systems. For instance, dynamic systems, such as the Earth orbiting the Sun, are subject to a *dynamic equilibrium*. However, knowing if this was a stable or unstable equilibrium was extremely difficult given the mathematical tools of the time. There are also systems with *multiple equilibrium* points, however this leads to the difficulty and sometimes impossibility of determining exactly which equilibrium a system will fall in to. Walras therefore decided that for every commodity there is exactly only one equilibrium point, and that this point is the price at which people are willing to trade at. He went on to develop a model, known today as a Walrasian model.

In this model, individuals are endowed with an amount of every possible good, and an individual's preference, or utility, towards goods is heterogeneous. The idea is that if people want to trade, the system is out of equilibrium and therefore a more optimal allocation of goods exists. If you could establish prices to trade at, people could trade to move to a more satisfied state. Once everyone was as satisfied as possible, no one else would want to trade and hence equilibrium is reached. Prices would be set by an auctioneer using one of the goods as money. As in a usual auction house, if there was more demand than supply, prices would increase and vice versa. This would be done across all goods, and once the price of every good was established, then and only then could people trade. His model assumed that people would act rationally and in their own self interest. Additionally, and importantly, the existence of the auctioneer, makes it mathematically simpler but also a *centralised* (thus unrealistic) system.[18]

Pareto was also hugely influential. He deduced that people would not make trades that would make them worse off, and hence every trade makes society as a whole better off. Today this is called *Pareto Superior* trading, and means people will only engage in trades that are good for both of them, which is what the model of trade employs in this report. With this, the market would eventually reach an equilibrium which is *Pareto Optimal*, where no individual can trade without making someone worse off, in other words, there are no more Pareto Superior trades to be made.[27, 1]

Both Pareto's and Walras' ideas were extremely successful, although Walras' model assumed certainty, or perfect information. It was Arrow and Debreu[27] who tied the theories of Walras and Pareto together, resulting in the *Neoclassical General Theory of Equilibrium.* They conjectured that markets will coordinate themselves and reach Pareto Optimal prices, however, this would be so regardless of uncertainty.

Now the basics of neoclassical economic theory have been laid down, we are in a better position to understand the fundamental limitations of this approach. One of the difficulties of economics as we know it in text books, is that it is very much focussed on macro results, and lacks the ability to understand the necessary micro interactions that can result in these macro trends. It is fair to say that this top down approach of analysis adopted by economics works very well for many fields, However, it is also the case that many of the biggest problems in nature are complex systems. These sort of systems benefit far more from a bottoms up approach. The Santa Fe Institute, in New Mexico, USA, does huge research in to complex systems. They are also extremely pro cross silos approaches to problems. They recognise and value the fact that many problems are better understood through utilising knowledge from many different areas of expertise. Economics and the Santa Fe Institute crossed paths in 1987 when Citicorp agreed to fund cross silos research into the field.[19] It was time to accept that economics may not be fulfilling its potential as a science, and was also time to utilise the expertise of leading researchers in various fields.

2.1.2 Economics as a Complex Adaptive System

At the Santa Fe Institute, days of discussion amongst economists, physicists, biologists, mathematicians and computer scientists regarding the current approach of economics lead to some interesting conclusions. It was realised that economics was fairly "behind the times" when it came to the mathematics it was based on. However, the scientists were somewhat impressed that economists had managed to bend the tools at their disposal so well creating such impressive results for such a long period of time.

A fundamental obstacle that the scientists and mathematicians saw was the assumptions in economic theory. It is true to say that assumptions are both apparent and useful across all disciplines however they must be correctly formed. Assumptions must reflect a simplified reality, they must not be a direct contradiction or an unattainable state. These concerns did exist in the days of Walras, however the progression Neoclassical theory continued with economists largely ignoring these concerns. At this time, people argued that so long as the output of the equations was realistic then the assumptions it was grounded on are of arbitrary concern. However, later on, people began to take a different turn. Herbert Simon from Carnegie Mellon University pointed out that the point of science was to explain, not predict. It does not suffice to validate a theory through verifying the end point or conclusion is as expected, it is necessary to validate the entirety of the theory - including the assumptions. [27, 19]

Probably the most debated assumption of economic theory is the model used for human behaviour. This is manifested in the assumption of perfect rationality which makes unattainable and unrealistic assumptions about both the world we live in and the way in which we engage with our environment. For instance, it assumes that we will always act in our own economic self interest, and that in making basic everyday decisions we do so with perfect information. For instance, when I make a decision to make a purchase of some arbitrary good, I have taken into account every other possible thing this could be spent on, assure myself that the increase in utility gained from purchasing this good is that which maximises my utility. I am certain that this specific Tesco offers the best price; not to mention accounting for whether or not this money should be put in to a savings account instead, which assumes knowledge of current and estimated future interest rates, government expenditure and so on. This decision making process evidently would result in a monstrosity of calculations on my part, involving information that is difficult to get, if possible at all. The truth however, is that this does not correctly reflect a model of human behaviour. In reality, we are not good at difficult calculations over split seconds. Our intelligence is in the ability to make such quick decisions with ambiguous and incomplete information on a regular basis. People also have the ability to learn and apply knowledge to new situations through pattern recognition - a very different picture than that of neoclassical theory. This perfect information is an example of something that will be explored in the model implemented - partly to see if it's necessary to have perfect information to reach an equilibrium, and partly to compare what happens with and without it; perhaps leading to gaining an insight into which offers a better reflection of reality.

In addition, the world we live in is often not simplified by neoclassical theory but altered. For instance, the existence of an auctioneer is something that (with the exception of an auction house!) never happens. It is not the case that the supermarket is centralised - they don't have to hold auctions for us to buy our weekly shopping. On the contrary, our networks are very much decentralised, and as such my model will veer from the Walrasian model by removing the auctioneer.

Another issue is that of the representation of time in neoclassical theory. In reality trawling for information, making trades and learning all take time whereas it's usually instantaneous in neoclassical theory. The importance of this is that to properly understand systems behaviour it is necessary to have an idea about the sort of time scales you are looking at.

A Step in the Wrong Direction

Again it is important to clarify that the neoclassical economic model is not entirely crazy. It is not true that the model of supply and demand moving to form an equilibrium is misclassified, it is a fair estimate, but rarely happens. It is also true that prices do sometimes converge, and markets can act as if they are in a form of equilibrium. However, the fact that these models aren't as bullet proof as something like "speed equals distance over time" leads us to some interesting questions. Namely, why is it the case that these models aren't often realised?

Is it really because of "exogenous shocks" to the economy constantly moving it out of equilibrium? Or are these "exogenous" factors actually "endogenous" to the system itself? Is the scope of the economic system wider and far more complex than originally anticipated?

In an attempt to answer these question it was necessary to discover exactly where these models came from. The physics that was borrowed by Walras and his compatriots to bring mathematics and economics together was from classical thermodynamics. However, at this point in time only the first law of thermodynamics was in existence.[27]

Thermodynamics studies the behaviour of energy flow in natural systems. From the study of this area, some fundamental laws have been observed in our universe. The *first law of thermodynamics* otherwise known as the *Law of Conservation of Energy*, says:

> *The change in a system's internal energy is equal to the difference between heat added to the system from its surroundings and work done by the system on its surroundings.*

All this means is that heat added to the system can only do two things - change the internal energy of the system or cause the system to do work. Mathematically this law is:

$$U_2 - U_1 = \delta U = Q - W$$

$$Q = \delta U + W$$

where U_1 is the initial internal energy, U_2 the final internal energy, Q is the heat transferred into (or out of) the system, W is the work performed by the system or on the system and δU is the change in energy.[20]

A consequence of this law is that if the energy of a system is conserved, the system is guaranteed to reach an equilibrium. The only thing that can move it from this equilibrium is adding energy from outside the system - exogenously. It states that the total energy of the system and its surroundings remains constant. Trivially this relates to the theory of supply and demand - prices will reach equilibrium through the tensions between supply and demand unless an outside force shifts it from this equilibrium. It also implies, that if the economy follows this law, wealth can not be created - the economy must begin with a finite set of resources that create a finite set of goods.

However, this leads to some debate, especially when you do look at the second law of thermodynamics. The second law of thermodynamics deals with not how to do something but instead constrains what can be done:

> *It is impossible for a process to have as its sole result the transfer of heat from a cooler body to a hotter one.*

It says we are restricted by nature to achieve certain kinds of outcome without putting a lot of work into it. Hence it is closely tied to the conservation of energy, just as the first law is. This law also has a lot of applications outside of physics since it is closely tied in to the notion of entropy. In terms of entropy, the second law reads:

> *In any closed system, the entropy of the system will either remain constant or increase.*

This means a system that goes through a thermodynamic process can never be returned to the exact same state it was in before, a definition used for the *arrow of time* as entropy of the universe will always increase over time according to the second law. Entropy is a measure of disorder, and hence the universe's disorder is always increasing and work has to be done in order to bring order to the system; the system decays into disorder and eventually comes to rest.[20]

These two laws lead us to a distinction between open and closed systems. A closed system is one which can exchange heat and work, but not matter with its surroundings. In contrast, an open system is a system where matter can also be added or removed from it, hence open systems interact with their environment. The question now is what does this all have to do with the economy? The economy is made up of energy, matter and information and hence is not an abstract notion - it exists physically and thus is exposed to the laws of physics. Energy enters the economy, we fight entropy as we fight to create order, and export disorder, hence obeying the second law through throwing out waste such as pollution and greenhouse gases. Thus economies are literally open systems. Since the neoclassical model is so heavily tied to the First Law of Thermodynamics, perhaps this is what has restricted the potential of economics. The conjecture however that the economy is in fact a complex adaptive system necessitates defining exactly what this means, which is the topic of the next section.

2.2 Complex Adaptive Systems

In Eric Beinhocker's book he argues that the economy is in fact a complex adaptive system and thus should be studied as such. In this section I will briefly describe what these systems are and attempt to explain how this relates to the economy.

Complex Adaptive Systems are a special case of Complex Systems, so we will begin by looking at what a complex system is. Complex systems are only recently becoming better understood, however the notion of complex systems has existed for over 100 years. In 1887, Oscar II, who was the King of Norway and Sweden, offered a prize for anyone who could tell him whether the solar system was stable. It was Poincare who showed that it was impossible to find a solution to the trajectory of just three planets interacting in a non-linear fashion. Fortunately he won the prize and this problem is known today as the *3 Body Problem.*[6] His findings were put aside for many years, and it is only relatively recently with the introduction of computerised simulations that people realised he had predicted chaotic motions and complex systems.

The definition of a complex system comes in many forms, here are just but a few.[6]

> *...you generally find that the basic components and the basic laws are quite simple; the complexity arises because you have a great many of these simple components interacting simultaneously. The complexity is actually in the organization the myriad possible ways that the components of the system can interact.* (Stephen Wolfram, quoted in Waldrop, 1993)

> *...to understand the behaviour of a complex system, we must understand not only the behaviour of the parts but how they act together to form the whole.* (Bar-Yam, 1997)

> *A complex system is a system for which it is difficult, if not impossible to restrict its description to a limited number of parameters or characterising variables without losing its essential global functional properties.* (Pavard, 2000)

Hence a complex adaptive system is a complex system with the added capability of learning and changing over time. Complex systems in general have many properties, such as emergence, short range non linear relationships, non-determinism, limited functional decomposability and distributed character of information. It is also important to note that a key focus of complex systems is on the interactions between elements of a system as opposed to the elements themselves. First I will explain these properties in more detail, and then go on to explain precisely what is the adaptive part of complex a*adaptive* systems.

Complex Systems

Non-determinism stems from the number of interactors and hence interactions in the system. Given that these agents interact in a non-linear fashion, we can see how chaotic behaviour can quickly emerge in a complex system and hence see how complex systems are intrinsicly non-deterministic. As an example of how chaotic behaviour can emerge, consider the quadratic map equation below.[27]

$$C_{n+1} = aC_n(1 - C_n)$$

If we vary the constant a we can radically alter the models behaviour. For instance, set C_0 = *0.1*, and a = *1.5*. C_n tends to 1/3, and then stays there forever, as shown by the cobweb diagram. This is an illustration of dynamical systems equilibrium - a fixed point attractor, so called as the equation is pulled towards this single point in the cobweb diagram. The cobweb diagram plots the value of the function at n on the x-axis and at *n+1* on the y-axis.

If however, we set a = *3.3* we get regular oscillations, which is known as a periodic limit cycle. Moving a more now, such that a = *3.52* we encounter a

Figure 2.1: Fixed Point Graph

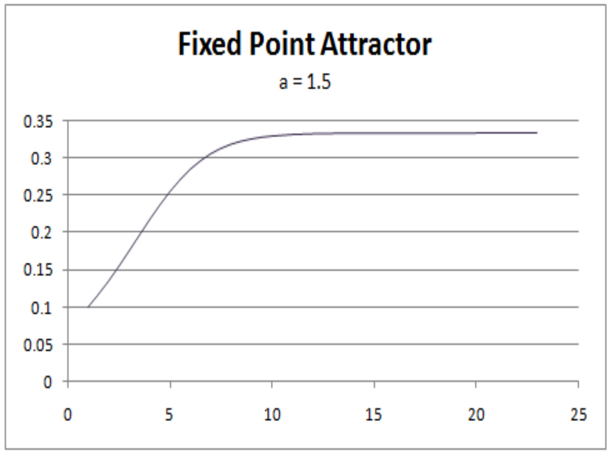

Figure 2.2: A cobweb diagram illustrating fixed point attractor

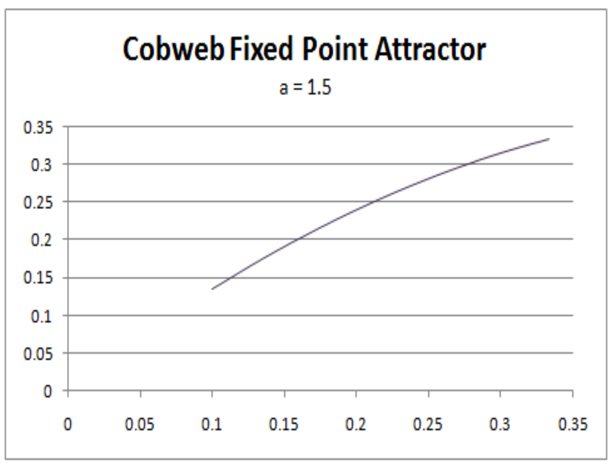

Figure 2.3: Periodic Limit Cycle Graph

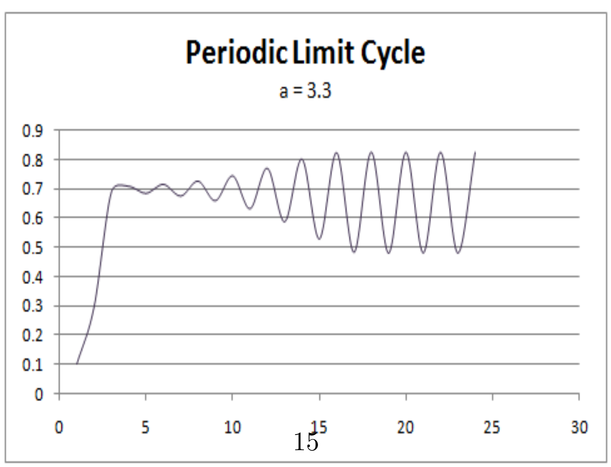

Figure 2.4: A cobweb diagram illustrating periodic limit cycle

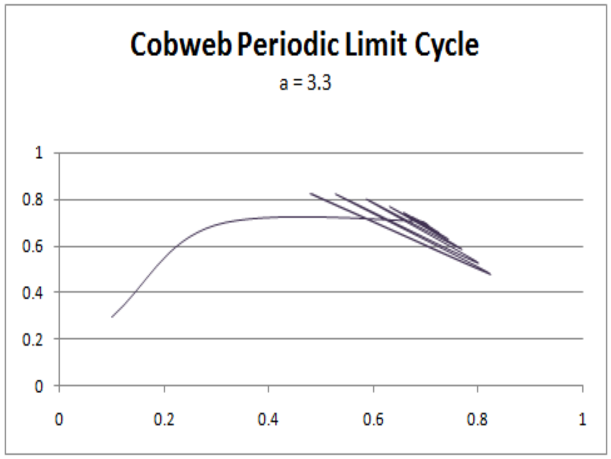

Figure 2.5: Quasi-periodic Limit Cycle Graph

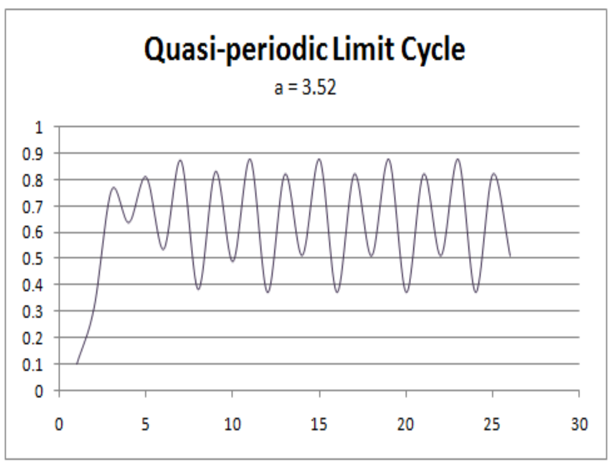

Figure 2.6: A cobweb diagram illustrating chaotic behaviour of quasi-periodic Limit Cycle

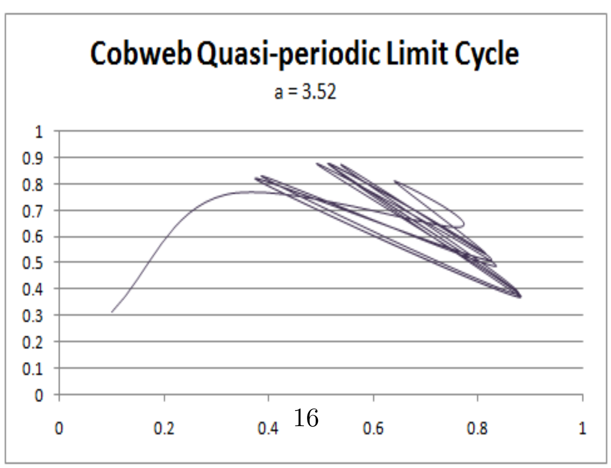

more complex pattern - oscillations within oscillations, known as a quasi periodic limit cycle. Finally set $a = 4$ and we reach *chaos*. This will never actually repeat itself, however, the system is bounded - it will never move out of the range 0 to 1.

As shown, these systems are very sensitive to initial conditions, and in addition are path dependent (the previous state is needed to calculate the next state). These two factors make these systems difficult or sometimes impossible to predict, even if the initial conditions are exactly known.

In addition complex systems often exhibit emergent behaviour - that is to say they may have properties that can only be studied at a higher level - the system is greater than the sum of its parts.[20]

Another interesting feature of complex systems is that they contain feedback loops. This is an example easily related to the economy. Firstly a feedback loop can either be positive or negative. It is linked to path dependence in that it is when something of the past effects something in the present. When the event is part of a cause-effect chain which forms a loop it is said to be fed back in to itself.[20] An example of this in the economy would be bull markets; when prices are rising, people believe that price rises are probable and therefore have an incentive to buy - an example of positive feedback.

The economy is definitely a dynamic system in that it changes over time. It also exhibits non-linearity in many areas from unemployment figures to the rate of technological advancement. The dynamics of the economy result from non-linear interactions between billions of individuals. The behaviour of the economy is unforecastable except in the very short term and hence we are unable to accurately predict its evolution - the economy is a complex system. Now we are convinced that the economy is a complex system, we can ask is it also adaptive? What is the difference between a complex system and a complex *adaptive* system?

A complex adaptive system is a complex system with added capability of being being able to change and learn from experience. It is the ability of systems to constantly react and adapt to changes in their environment, or to changes in the way in which other agents behave.

The economy is both a complex system and adaptive. Billions of people interact, learn, communicate and generate new strategies, business models and technologies everyday. The economy is constantly evolving, people cooperate to achieve goals, share knowledge and this experience, learning and application of new knowledge is what distinguishes a complex adaptive system from a complex system.

Evolution is also an important notion in this new perspective of the economy as a complex adaptive system. The evolutionary process is responsible for enabling new discoveries contributing to growth in order and in complexity.

2.3 Evolution

Evolution is an algorithm that searches some space for fit designs. For evolution to work, certain criteria must be met. There must exist a *design space* in which all possible designs are contained. These designs must be able to be encoded into a *schema* and a *schema reader* must exist to decode these designs into *interactors* (these readers may be endogenous to the interactors). Constraints in the environment in which the interactors live form a *fitness function*, rendering some interactors fitter than others - the criteria for selection. The algorithm of evolution consists of only three stages - *differentiate, select, amplify.* Differentiation is rendering different schema into interactors in the environment, and tweaking interactors. Selection is the selection of fit interactors in the environment according to the fitness function for amplification. Amplification is the spread of good designs in the physical environment, making some small changes, and the occasional big one. In the economy you could hypothesize that business models are what are differentiated. The market is the selector - bad business plans don't survive in the competitive market, and amplification is the spread of knowledge and imitation of good business plans.

Evolution makes many small alterations and some big ones to the interactors in the given environment which efficiently searches a massive design space, and since the selection picks the strong ones, the designs evolve and improve over time. The economy is therefore adaptive because we learn and experiment with what we have, searching for most profitable ways of doing things.

It is probably easier to see now why it is so difficult to forecast the economy over anything but the extremely short term. Sensitivity to initial conditions, dynamic complexity and path dependence all add to this difficulty. However, as briefly mentioned, these sort of problems may be better studied using a bottom up approach which is where Agent Based Modelling comes in.

2.3.1 Genetic Algorithms

2.4 Agent Based Modelling

Agent Based Modelling provides a way of growing artificial societies of sorts. A key point of the previous section was that although the economy may be unpredictable and unforecastable, that doesn't mean economics is a lost cause. Science is not about simply predicting, but explaining how certain phenomena emerge from micro interactions. This is exactly what Agent Based Modelling allows us to do.

> *.. one must show how a population of boundedly rational (cognitively plausible) agents, interacting locally in some space, could actually arrive at the pattern on time scales of interest - be it in wealth distribution, spatial settlement pattern, or pattern of violence. Hence to explain macroscopic social patterns, we try to grow them in multi agent models.* (Joshua M Epstein, Generative Social Science)[24]

Agent based modelling involves creating autonomous agents in some space, in this case an economy, who are heterogeneous and interact with each other in a decentralised fashion (i.e. through local interactions). In my model agents will be heterogeneous in how much they can produce of the two goods and later on also in how much they can remember. By running the simulation I will be able to observe and analyse any emergent behaviour. It is also important to note the micro specifications that result in particular macro phenomena are not definite solutions, moreover they are candidate explanations that prompt further investigation, especially if there exist multiple micro specifications that result in a macro structure. In this case it would be necessary to do further work in order to determine which is the most likely candidate solution. Agent based modelling is also a tool to subject theories to stress testing. In the context of my project, relaxing assumptions of neoclassical economic theory could help deduce whether or not these assumptions are necessary to produce a specific macro phenomenon.

In the quotation above, it was noted that agents have bounded rationality. Their rationality is bounded in two ways. Firstly through information. Agents do not have access to global information (although it is possible to create this as a network in order to observe any differences it creates). Secondly they are bounded by computational power, in that it is not infinite.

Many agent based models have proved to be a huge success and the use of such models is rapidly gaining pace as a way of studying such complex systems. A good example would be the implementation of a simulation of the Anasazi society by Dean, Gumerman, Epstein, Axtell, Swedlund, McCarroll and Parker. In this simulation they attempted to grow a "500 year spatio-temporal demographic history - the population time series and spatial settlement dynamics of the Anasazi,"[24] which was tested against empirical data.

This was a society that existed in a valley in Arizona between 800 and 1300AD, but then vanished. The study managed to conclude that it could not be environmental factors alone that resulted in the demise of the Anasazi - a huge step forwards for the long search in to what had happened. [23]

2.5 Inspiration

After reading the fantastic paper written by Allen Wilhite titled "Bilateral Trade and Small-world Networks"[1] I developed a strong interest in his conclusions and felt that his model would be a good place to start. Wilhite implemented an agent based simulation in which agents were able to produce or trade one of two durable goods. His aim was to explore the efficiency of various network topologies with respect to search, negotiation and exchange. He also experimented to see the effect that the various network topologies [1] (which will be defined later) had on issues such as the speed and extent of price convergence. I think his model could be a solid foundation to build on and in addition I

[1] The topologies differed in who each agent could trade with. For instance one topology had agents in disjoint groups and in another every agent could trade with every other agent.

feel that the report was written well enough for me to attempt to replicate his work. Therefore, in the next section I will start with my initial model which is that of Wilhite's paper. Wilhite went on to write a second paper in 2003 titled "Self-organising Production and Exchange"[22]. In this paper he introduced the notion of transaction costs to reflect the cost of shipping and other expenses incurred in trade. I felt that this was a realistic addition and therefore have also incorporated this into my initial model.

Wilhite's work however is not the only work to have had such a strong influence on my project. In one of the first works at Imperial College London in the area of complex systems and social dynamics, Kelvin Au did his individual project in 2005. He implemented a simulation titled "Dynamics of Human Behaviour: Evolution of Hierarchical Groups"[7]. His report gave extremely sound explanations on complex systems in a very accessible way and the base of my knowledge in this area did indeed came from his report. He also had a lot of work on various frameworks available for simulations and although in the end I decided to implement it myself it was good to have all of this research integrated in one document. He also shared the notion of interactions with Guèrillot.

Camille Guèrillot wrote an MSC report and completed the implementation of an agent-based simulation at Imperial College London in 2005. He wrote a simulation on the "Dynamics of Human Behaviour"[6] which gave me a great insight in to the complexity of the field. He utilised Au's technique of using interactions which has inspired my work. In his model every agent got an opportunity to perform an interaction on every iteration. In an interaction, an agent would engage with a neighbouring agent, resulting in one of several possible actions. [2]. This notion is reflected in my model described since extensions will be modelled as interactions - to start with there is just produce and trade, but later learning can be seen to be an interaction, as can reproduction.

Jie Shen completed an independent study option at Imperial titled "Dynamics of Human Society: Introduction to Multi-Agent System Based Research in Social Sciences".[3] His account provided a lot of information on multi-agent systems and the research being carried out in the area which definitely helped further my understanding.

2.6 Related Work

Aside from Wilhite's model, another interesting and notable work was "ASPEN". This was the implementation of an agent based model designed to simulate the ASPEN economy, documented in the paper "ASPEN: A microsimulation of a model economy"[4]. This employs a Monte-Carlo simulation, in which agents are designed to be "real-life-economic-decision-makers". In the world, households exist who are employed, or alternatively on social security benefits and strive to earn an income to consume goods or save money. Multiple industries are modeled through creating agents as firms, and these firms set prices using a genetic algorithm learning classifier system permitting the development

[2]To name a few interactions: Talk, fight, flirt,

of pricing strategy. In addition, the economy is governed by a single agent, and a financial sector is modelled. The simulation produced dramatic results, including the emergence of business cycles. It's aim is to be improved and enriched enough for it to become useful as a forecasting tool.

The simulation documented in this report is far more abstract than that of the economy of ASPEN. However, the aim is not to create a model used in forecasting the economy of a small state, moreover, it is an investigation into the potential of agent-based models as a tool for gaining insight into how trends witnessed in today's world might emerge. In fact, the abstraction actually still allows for some extremely interesting results, emphasising that even the simplest of models have their contributions.

Chapter 3

Model

3.1 Introduction

In the background section we established that perhaps a better way to study economics is through a bottom up approach, something to which agent-based modelling is well suited. Although the simulation to be implemented is extremely simple in that there are only two goods and there is no distinction between individuals and firms, it is hoped that it will allow for some interesting analysis of the assumptions made by Neoclassical theory in terms of their realism and necessity, and give a realistic representation of wealth distribution. In addition, employing evolution will allow an insight into the importance of various attributes of agents in different contexts with respect to the model at hand.

The model to be implemented is that of a basic economy in the form of an agent-based simulation. In the artificial world, autonomous agents are able both to produce and trade one of two durable goods, Good 1 and Good 2. The simulation is a series of ticks, and on each tick every agent is given the opportunity to perform one of two actions. An agent is constrained either to producing one of the two goods, or exchanging one good for another, and each agent's action is carried out sequentially. The act of exchange depicts bilateral trade in that agents swap goods for goods. In the model, Good 2 is infinitely divisible, acting as money, whereas Good 1 can only be traded in whole units.

Initially the model is simple and agents act in a rational, strategic, myopic manner insofar as they have a common goal of maximising their utility and do not try to trick other agents by misleading them. However, as the simulation is enriched, other aspects such as evolution, learning and a notion of memory will be brought in. These will be covered in due course but for simplicity, there will first be an explanation of how the initial model has been implemented and an evaluation of the findings. From here we will be better equipped to understand the natural progression and direction of the focus.

3.2 Initial Model

In 2003, Allen Wilhite[22] implemented a simulation of bilateral trade between agents in a produce - exchange economy. The initial model is an implementation of this work since it will provide a solid foundation upon which to build. From here, extensions will be tried and evaluated, as well as the initial set ups in order to better understand how the micro interactions of agents may lead to macro trends. The project, therefore is somewhat experimental and hence abstraction and dynamism are essential. Attaining technological goals will be discussed after the model is explained.

First, on each tick or iteration of the simulation, every agent has the opportunity to produce or trade. An agent is allowed to produce only one of the two goods *or, but not also,* trade with another agent on each of his turns, although he may be picked as a trade partner by another agent. An agent's decision is deduced by calculating which action will maximize their utility.

Utility was first introduced by the mathematician Bernoulli; however, the importance of this revelation went unnoticed for many years. It was only 60 years later that an English philosopher, Jeremy Bentham,[10] independently discovered this notion. He proposed that pursuing your own best interest did indeed translate into making economic decisions. He went on to introduce the measure of pleasure or pain to be one's utility, a measure in *utils.* Furthermore, one would make economic decisions based on maximising one's utility. It is important to note however, while Economic theory today tends to regard utility as an abstraction from pleasure and pain, it is rather an order of preference with no link to explaining mental processes from which it stems and, in addition is only a relative measure.[8] Utility of agent i is calculated in the model according to the Cobb-Douglas Utility Function:

$$U^i = g_1^i g_2^i, \quad i \in \{1, ..., n\}$$

where g_1 and g_2 are the amount of Goods 1 and Goods 2 possessed by agent i respectively and n is the total number of agents.[1]

It is worth noting that this is a symmetric utility function, in that an agent has no innate preference towards either good. This means that the two agents' desire for the goods is inherently equal, but deals are available when the differences between amounts of both goods are apparent across the two agents.

Having established the goal of an agent, let us move on to the next aspect - production. For production, an agent has a simple unique production function for each good. An agent may produce r^i of g_1 and s^i of g_2. Formally,

$$\Delta g_1 = r^i; \quad \Delta g_2 = s^i \quad r, s \in \{1, ..., k\}; \quad i \in \{1, ..., n\},$$

where r and s are randomly determined integers at initialisation lying in the region 1 to k, and n is the number of agents.[1]
The constant k in Wilhite's model[1] was $k = 30$, and this value will also be used here (although changing it will be possible).

For trade, otherwise known as exchange, there are three stages: *Search, Negotiation,* and *Exchange.* Search is the act of agent i finding a partner with whome to trade. At first, m agents will be randomly selected from the set of agents with whom agent i is allowed to trade. Agent i will calculate his *Marginal Rate of Substitution* or *MRS* as well as the *MRS* of all of the m agents.

The *Marginal Rate of Substitution* is the amount of Good 2 an agent is willing to give up for another unit of Good 1. Using the utility function, the *MRS* is given by [1]

$$mrs^i = \frac{U_1}{U_2} = \frac{g_2^i}{g_1^i}$$

where

$$U = U(g_1, g_2), \quad U_1(g_1, g_2) = \frac{\delta U}{\delta g_1}(g_1, g_2), \quad U_2(g_1, g_2) = \frac{\delta U}{\delta g_2}(g_1, g_2) \quad i \in \{1, ..., n\}$$

A difference between the MRS of two agents is indicative of an opportunity for mutually beneficial exchange. If this is the case between the searching and selecting agent, we move on to negotiation. Negotiation is the act of deciding a price at which to exchange goods and of deciding on the amount to exchange. The trading price between agents i and j is calculated using:[1]

$$p_{i,j} = \frac{g_2^i + g_2^j}{g_1^i + g_1^j}, \quad i, j \in 1, ..., n$$

Since this price is per unit, the agents will partake in hypothetical trading to decide the quantity to trade and this ceases when it is no longer increasing the utility of either of the two agents. Hypothetical trading is necessary since you want to search through the m agents for the best deal and hence a trade can't execute it until all m agents have been reviewed.

In summary, the algorithm of the initial model on each iteration is given below.

1. Calculate the change in utility you would gain from choosing to produce Good 1 and commit it to memory.

2. Calculate the change in utility you would gain from choosing to produce Good 2 and commit it to memory.

3. Calculate your Marginal Rate of Substitution, MRS.

4. From the potential agents you can trade with, select m and for each agent:

   (a) Look at the agent's MRS, if yours and theirs do not differ, disregard the agent and move on to the next one. If they do differ, continue with the steps below.

(b) Calculate the price at which you will trade.

(c) Engage in hypothetical trading until either of your utilities is no longer increased.

(d) Remember the hypothetical trade and return to step (a).

5. Compare all items in your memory and choose the action which yields the greatest increase in utility.

6. Execute this action.

3.2.1 With whom can an agent trade?

In the initial model, who an agent can trade with is imposed purely by the system. However, it would be interesting to allow for the evolution of trade networks through allowing agents to expand their network of contacts over time. There are several key contributors to the reason for initially restricting an agent's possible trade partners. One is the simplicity of the implementation, making it easier to verify the correctness of the basic algorithm through employing minimal complexity. Secondly, it would be interesting to evaluate the evolution of trade networks with respect to the initial network employed. Finally, in order to properly evaluate the effect of learning on the running of the simulation, it is necessary to have a method of comparison - to know the difference with and without this feature. This will allow sounder conclusions to be drawn from any changes in macro trends or micro behavior. There will be four networks to choose from, depicted in the Figures 3.1, 3.2, 3.3, 3.4.

For simplicity, imagine the economy as a set of agents organised as a ring lattice around the edge of a circle. In a Global Network, Figure 3.1 every agent can trade with every other agent.[1] In the Local Disconnected Network, Figure 3.2 , each agent is part of a distinct subset of agents, called *districts*. The subsets are both disjoint as an agent can only appear in one set, and exhaustive in the sense that every agent belongs to a set. This network is analogous to an *autarky*. This simply means it does not take part in "international" (cross district) trade - it is a *closed economy*. Finally, the Local Connected Network, Figure 3.3 and the Small-world Network, Figure 3.4 are crosses between the Global and Local Disconnected networks. The number of agents together with the number of sets will be configurable. It is also worth noting that the m agents selected as potential trade partners will be randomly selected in the initial model since there is no notion of memory.

Having established an understanding of what interactions occur between agents and the various contexts in which they do so, let us now evaluate the findings of the initial model. The evaluation is based on the running of 10 simulations (of each set-up), varying the initial conditions. Upon completion of a simulation, a PDF document is generated. This contains data on macro trends, specialisations, and the strategies of particular agents. It is the comparison of these outputs that is the basis for conclusion on the legitimacy of the model with

Figure 3.1: Global Network

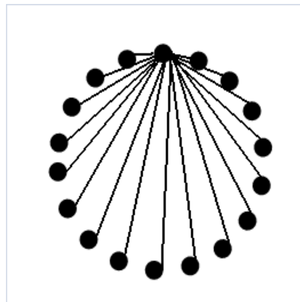

Figure 3.2: Local Disconnected Network

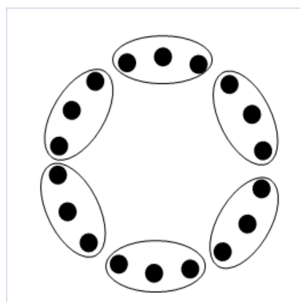

Figure 3.3: Local Connected Network

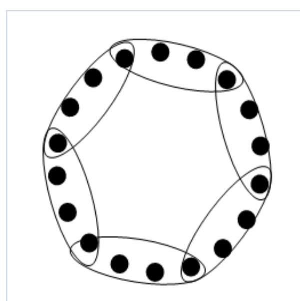

Figure 3.4: Small-world Network

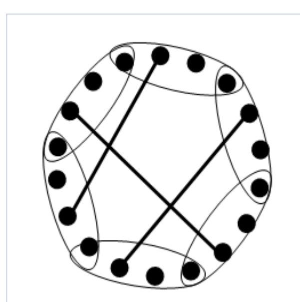

respect to the real world, gaining insight into the impact of initial conditions on the evolution of the simulation, and any limitations that the model presents.

3.3 Evaluation

There are several key questions that shall be addressed throughout this section - namely:

- Is there price convergence?
- Is there a difference in the dispersion of prices for the different initial network topologies?
- Is there specialisation in what an agent chooses to produce (or purchase) or specialisation concerning with whom they choose to trade?
- How do topologies effect the distribution of wealth in societies and is this distribution similar to that of the real world?

3.3.1 Prices

In every round, every agent has the opportunity to search through all agents in their district in order to locate a trade partner. As such, the topologies lie on a continuum between two extremes. On the one hand, in the Global Network every agent can trade with any other agent, and on the other, in the Local Disconnected Network agents are constrained to being able to search through only a subset of the entire population of agents. As a result, from the simulations run, it was clear that the dispersion of prices differed across network topologies. Prices were measured in terms of the amount of Good 2 you would give up for one unit of Good 1, thus the price refers to the price of Good 1. In the Global Network, dispersion, measured as the standard deviation of the average price, was fairly low, on average the deviation was 0.024 (see Figure 3.6 for price over time). This is best explained by the fact that since every agent is able to search through the entire population to find an optimal trade partner, there are few trades that go unrealised. However, in the Local Disconnected network the opposite is true. The isolation of traders means that there are many opportunities for trade that are missed, and although each district converges to an average price with little deviation, the global standard deviation is approximately 3 times larger than that in the Global Network. A great illustration of just how much deviation from the global average price is possible, can be seen in Figure 3.5.

Let us now consider the topologies that lie in the middle. The Local Connected and Small World Network differ in that although trade occurs locally within districts, certain agents are made to be bridges between districts, or "crossover agents". In the Local Connected Network however, the crossover agents only overlap with a neighbouring district, and they also only overlap with one other district even if there are multiple crossovers. This means that

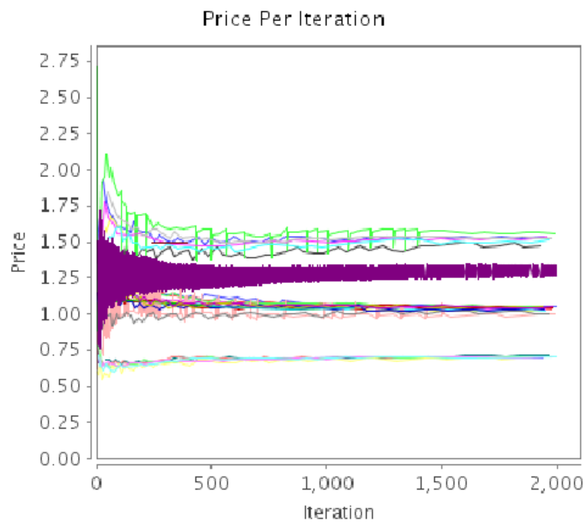

Figure 3.5: Illustration of price convergence in a Local Disconnected Network

although the majority of trade occurs locally between agents, goods can propagate through the network and spread globally. This gives lower search costs, since agents can only search through a subset of the population, but on the otherhand means that the average path length between agents is increased from the Global Network. This would suggest that not only convergence would be slower, but also the dispersion of prices should lie somewhere between the two extremes of the Global and Local Disconnected Network topologies. In fact, the difference between the deviation from the average price was only marginally different from that of the Global Network as shown in Table 3.1. In addition, the speed of convergence was also fairly close. As for the Small World Network, the dispersion was lower still (see Figure 3.7). It seems that the continuum of network topologies does not correlate linearly to the difference in prices across topologies. This shows that there is some efficiency in both the Local Connected and Small World Networks. Even though the path lengths for goods to reach an agent is increased, the cost of search is greatly reduced relative to the Global Network, and this reduction does not make a very significant difference on price dispersion. This gives rise to further investigation, and the topic of network efficiency will be covered in Chapter 6.

Contrary to the Neoclassical view of equilibrium, I found that prices did not converge to a single uniform price. However, the oscillation of prices did dampen considerably with time across the simulation. At first, prices (the amount of Good 2 an agent would be willing to give up for one unit of Good 1) fell within the average range of 0.5 and 1.5. By approximately 150 iterations, this fluctuation had reduced to between 0.85 to 1.1, and by 1000 iterations, it had reduced further to, on average between, 1.05 and 0.95. These were taken as

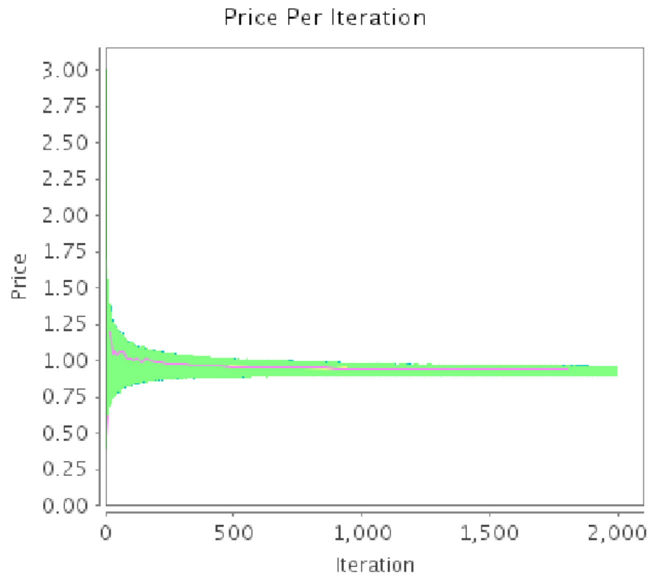

Figure 3.6: Illustration of price convergence in a Global Network

an average over 10 simulations of the Global Network with 400 agents, running for 2000 iterations. The Global Network was chosen since it best coincides with the Neoclassical assumption of perfect information - being able to search through every agent in the population.

Recall the Neoclassical view of price equilibrium was rooted in the first law of thermodynamics, stating that if energy is conserved, then the system is guaranteed to reach an equilibrium. It was thought prices would reach equilibrium unless an exogenous force was to shift it from this equilibrium. In the simulation, the reason for price fluctuation is entirely endogeneous to the system; it was the production of goods by actors in the economy that caused persistent price fluctuations.

Let us now address the reason for the continuous fluctuations in prices. It seems to stem from the possibility of production. Since agents often chose to produce, the stock of goods changes frequently which in itself prompts price adjustment. Although prices did fluctuate, it was always within a clearly defined range. This is common in the real world, especially with commodities where the amount being produced is not consistent. For example agricultural products often suffer potentially damaging price fluctuations. This can be due to poor crop yields stemming from weather, disease and so on, or to restricted supply, or to over production leading to insufficient demand for the quantity produced.

Full details of average prices and standard deviation, averaged over 10 simulations, each with 400 agents and 20 districts with only 2 crossovers (if appli-

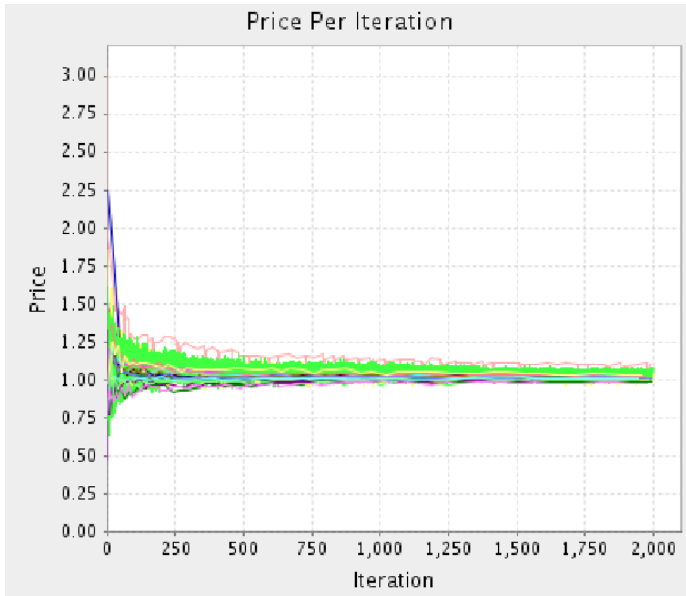

Figure 3.7: Illustration of price convergence in a Small World Network

cable) are given in Table 3.1.

Topology	Price	Standard Deviation
Local Disconnected	0.948	0.08
Local Connected	0.958	0.03
Small World	1.01	0.03
Global	0.935	0.024

Table 3.1: Table showing average and standard deviation of prices with different network topologies

3.3.2 Specialisation

Throughout the simulations that were run with the initial model, a common theme emerged. The ratio of production to trade was highly skewed towards production. This is shown by looking at the percentage of agents who specialise with respect to the interaction they perform most frequently - production or trade. In Wilhite's paper, [22] he categorised agents on a continuum. At one end of the spectrum are *pure producers*, agents who choose to produce at least 99% of the time. On the other end are *pure traders*. Similarly, these are agents who choose trade at least 99% of the time. In between lie the *heavy producers*

Figure 3.8: Continuum of specialisation with pure producers, PP, to pure trades, PT, with 0% to 100% trade

Topology	Pure Producer	Heavy Producer	Heavy Trader	Pure Trader
Local Disconnected	43.7%	55.3 %	1%	0%
Local Connected	43.4%	55.5 %	0.8%	0.3%
Small World	47.9%	50.6%	1.3%	0.2%
Global	48%	50%	1.4%	0.6%

Table 3.2: Percentage of agents in each category with respect to network topology

and *heavy traders*. These are agents who produce or trade, respectively, more than 50% but less than 99% of the time. An illustration of the continuum is shown in . The continuum can be viewed as illustrating extents of specialisation, with extremely high specialisation at either end, and little specialisation towards the centre.

Before discussing the degree of specialisation, if any, across the agent populations it is perhaps important to establish why this is of use. By seeing if agents do specialise, it is possible to learn about the model. The strategies that agents develop of their own accord in order to become optimal operators in the economy (with respect to their own potential and not a global optimal) become apparent. By examining what decisions are made on a micro level, it is easier to infer why these decisions are made, and hence better prepared to answer exactly why certain macro trends occur, whether or not the model is realistic and also to gain insight to any shortcomings of the model.

As may be expected given the above, the majority of agents fall in to the pure and heavy producer categories, shown in Table 3.2. In addition, it is often the case that agents with the highest utilities are also, more often than not, in these categories. Let us therefore consider what it is about production that is so much more appealing than trade. Why is it that, the majority of the time, even when an agent can search through all other agents, are there so few opportunities in which trade proves to be more beneficial than production? Is this really a realistic reflection of the world we live in?

In the model production involves no sacrifice. An agent simply adds to the stock pile of one of his goods - he give up nothing but time, the same amount of time another agent gives up for trade. With trade however, an agent gives away a stock of one good in exchange for a stock of the other. The benefit is only really presented when the quantity of goods you possess is heavily skewed, and since utility is calculated to be the multiplication of your stock of goods, agents are inclined to keep the stock of their goods similar to maximise utility. Put differently, the symmetry of the Cobb-Douglas utility function results in "balanced consumption" leading to greater utility.

The effect of topology on the percentage of agents falling into each category is fairly small. In Local Disconnected networks, the percentage of agents falling in to the pure and heavy traders category is reduced by approximately a quarter, illustrated in Table 3.2. This is most likely due to the fact that knowing fewer people reduces the probability of someone in your district being a suitable trade partner. What is interesting however, is the fact that there is so little difference between the number of traders in the Small World, Local Connected and Global Networks. This is illustrative of the fact that the links provided by the crossover agents in the Local and Small World Networks allows for flow of goods around the network, and are of little hindrance to the realisation of trading opportunities.

Also worthy of note is the relative rush of trades at the beginning, consistent across all network topologies. Figure 3.9 illustrates this rush, and, as shown, the number of trades falls fairly rapidly at the beginning (the range falls by an average 50% within the first 500 iterations [1]) and levels off for the duration of the simulation to a range of (on average) 10.7 to 5 trades per iteration. This is a significant fall and its explanation brings us back to Pareto. When the simulation begins, there are many trading opportunities due to the uniform distribution of goods - their endowments. Feldman showed when studying equilibrium characteristics of bilateral trade, [1] that as long as the agents are in possession of more than 0 of one of the commodities, that the pairwise optimal allocation is also a Pareto optimal allocation. That is to say that by selecting pairs of agents to trade, upon reaching a steady state, no more trades could occur that made both parties better off. However, in the model there is a twist. The optimal allocation of resources changes with time due to the possibility of production. Therefore, opportunities for trade are more prominent in the beginning, but as people exhaust those opportunities, the balance of goods and allocation of resources is optimised and there becomes little room for mutually beneficial exchange (as explained previously). However, it is important to note that the trades are not reduced to 0. This is due to the fact that trades are occasionally made more beneficial than production as the stock of goods of some agents become more skewed, since, as we know there are some pure traders in the economy. If we began with a finite set of resources, (if, for instance, production weren't possible) as the steady state is explained in Neoclassical economics, we would reach a point where no more trades were possible and this point would coincide with the point at which a price equilibrium is reached. This is not the case. Production means that thre are continuing trade opportunities that prove to be more beneficial that production (although they are, admittedly, few) and accounts for fluctuations in prices.

Let us consider the contributing factors to the emergence of this specialisation. What distinguishes agents who specialise in production and trade? The

[1] This is an average over 10 simulations of the Global Network.

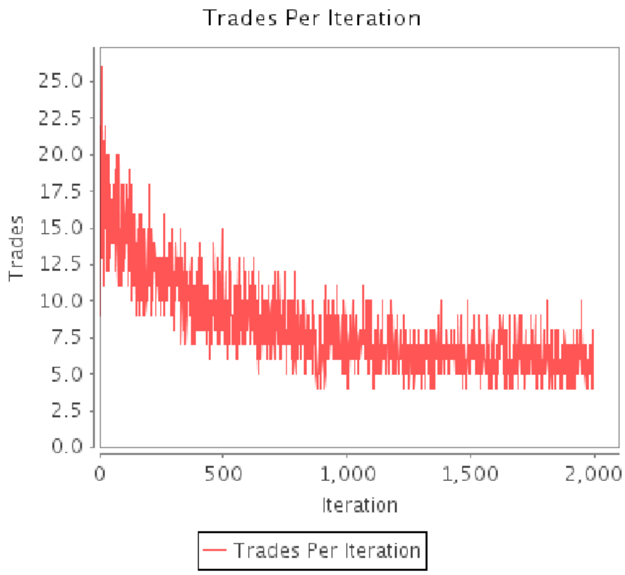

Figure 3.9: Trades over time for a Local Connected network

answer to this is the production functions of the agents. Agents develop repetitive strategies that they employ in the simulation based on their production functions that allows them to reap the maximum benefits of what they have been given. It is possible to profile the agents that fall in to these categories.

Producers

These are the agents who have high production functions. They are well equipped in the economy. Production virtually always offers better results than trade. They can be seen as the self sufficient sector of the society.

Often these agents simply alternate in the production of the two goods, keeping their stock piles close together and enjoying an easy life with high utility. The majority of the time, this is a one-period cycle - they produce Good 1 then Good 2 and so on. These agents often never initiate trade. Furthermore, they are not even picked by others as trade partners. An illustration of the movement of goods for a self sufficient agent is shown in Figure 3.10.

Self-sufficiency is not the only strategy of pure producers and also was shown not to be the best. Some agents were particularly proficient in production of only one of the goods. They specialised even further in that not only were they pure producers, but they also only produced one of the goods the majority of the time. On average, the good that they were most proficient in would be produced 97.5% of the time. What is interesting is the dependence of these agents on

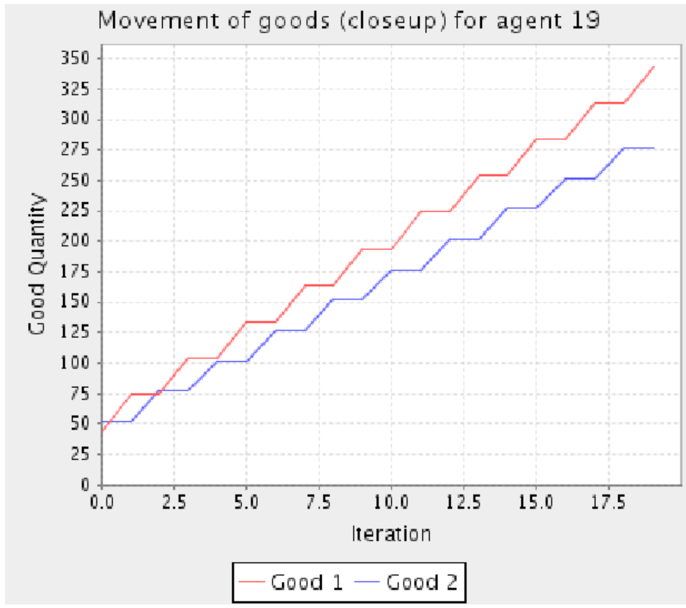

Figure 3.10: A close up of the movement of goods for a self sufficient agent, showing alternation between production of each good

the need of other agents to trade with them. They never initiated trade, but their position meant they were good candidates for many other agents who lacked production proficiency in the good they were producing. This strategy, on average, was apparent in approximately 22.7% of the pure producers. These producers were the consistently the most wealthy agents in the simulation. This is probably due to the fact that they achieved exchange without having to give up the opportunity to produce. Their balance of goods became closer because *others initiated trade with them.* In a sense, they could exchange for free because they didn't have to give up time.

The general profile for heavy producers was agents who still had skewed production functions but often neither production ability was particularly poor (ranging from 10 to 23 for those with the highest utility). These often produced the good they were proficient in for a number of rounds, and then trade for the other good before falling back to production. In addition, they had on average a third of the number of agents initiating trade with them compared to the pure producers who specialised in the production of one good.

Traders

Pure traders generally spent 100% of their time trading. They typically had extremely low ability in production (ranging from 1 to 7 for those with the highest utilities) and their wealth was unequivocally lower than the producers. Due to their poor production capability they often made margins through purchasing

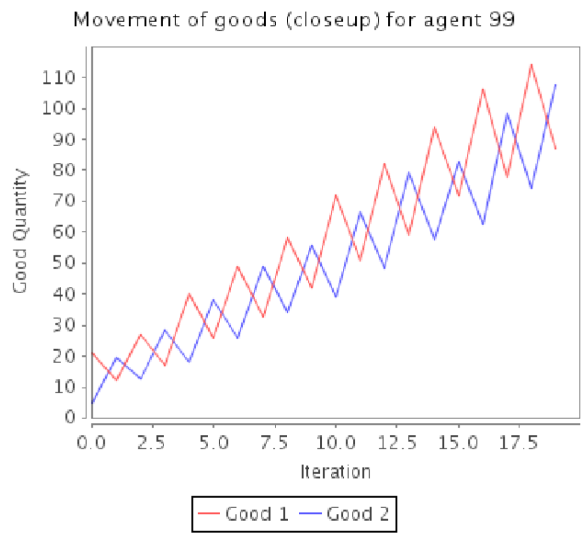

Figure 3.11: Illustration of the movement of goods for a pure trader

goods in one round and selling in the next. Although they frequently initiated trade, on average no agent initiated trade with them. Their lives were bleak, and they were the poorest of the agents, but trade was the only way of survival. Pure traders didn't specialise in the good they chose to buy. They, like the pure producers, switched between buying and selling Good 1 from round to round giving a 50/50 split. An illustration of the movement of goods for a pure trader is shown in Figure 3.11

However, they too specialised further not in what they traded but in with whom they traded relative to the heavy producers and traders. On average in the Global Network, out of 1000 trades, there would only be 66.2 distinct trade partners out of 400 possible trade partners. This shows that pure traders had "regular clients" in that they would repeatedly return to the same agents to trade.

Heavy traders also often suffered similar poor ability in production, although this was skewed and hence they, like heavy producers, would have one round of production and then several rounds of trade.

Effect of topology

Table 3.2 shows that although the Global Network has the most trades out of all the topologies, it actually also has the most pure producers. It seems that as the topology restricts the number of agents known, more agents migrate categories, from the two extremes of specialisation, pure producers and traders, to the middle of the continuum, heavy producers and heavy traders. It is easy to see there are fewer pure traders in the topologies other than the Global Network.

It is not feasible for agents to find a suitable trade partner (one which makes trade more profitable than production) in every round when they are only able to search through a small subset of the population. However, it is less easy to understand why the number of pure producers is considerably lower in the Local Connected and Disconnected Networks relative to the Global and Small World Networks. The most likely explanation is to do with competition in the market.

In the simulation the notion of competition relates to the fact that the more agents with whom one must compete for a trade partner, the more likely it is that the deal will be taken by another agent. This means that an agent who could have been a good trade partner now has a better aligned stock of goods for maximising their utility, so trading with certain formerly useful agents is now of no use to them. The Small World Network and Global Network offer the most competition in trades. In the global network, every agent searches through every other agent and therefore the chance of a trade being made that jeopardises another agent's opportunity to trade is high. In the Small World Network it is slightly more complicated. In Small World Networks a district can be connected to any other district, or more than one if there are multiple crossover agents. This means that competition across districts is higher as more crossover agents compete for the best trades. Therefore, crossover agents may actually have difficulty exploiting their position because often all the optimal trades are gone and therefore they are more likely to become pure producers. Notice that in the Local Connected Network there are far more heavy producers. It is important to be aware that this change is not purely the result of competition. It is a combination of competition allowing crossover producers to have occasional extra trade, and of the network being more restrictive, therefore encouraging pure traders and heavy traders to move to production - something not exclusive to crossover agents.

This has the knock on effect that since they are no longer bringing in good deals for local agents in their district, they too are more inclined to succumb to production. Therefore, open networks lead to more trade, but also to more extreme specialisations, since competition in the world increases. On the other hand, in the Local Connected and Disconnected networks, the lack of competition actually leads to less extreme specialisation since opportunities are available more of the time for some agents, although globally it leads to fewer trading opportunities. This is an interesting observation as it is illustrative of the connections between macro and micro trends which may seem superficially to be in opposition to one another but which are not in fact mutually contradictory and between which there are elements of causality and symbiosis.

3.3.3 Wealth

Wealth and its distribution are important indicators of the extent to which reality is captured in the model. This can also be related to countries of todays world, and serves as an indication of allocation of resources, poverty, flaws in policy and so on. Its importance is a reason to investigate how this evolves through the simulation and why.

Let us first establish what wealth is in the simulation. Wealth is measured as the value of assets minus the value of liabilities. Since in this model liabilities, or debt, do not exist, wealth is simply the value of the stock of goods held at one time. This leads to the need for a definition of what is meant by value in the context of the model. Since prices are measured in terms of how much Good 2 an agent would be prepared to give up or *pay* for one unit of Good 1, it is fair to say that the value of an agents assets is:

$$w_t^i = p \times g_1 + g_2$$

where w^i is the wealth of *agent i* at time *t*   *g1, g2* is the stock of Good 1, or Good 2 respectively for *agent i* at time *t*,   *p* is the average price of Good 1 at time empht

Having defined wealth, let us now identify what and how to evaluate the distribution of this wealth. Firstly, it would be interesting to see whether the global distribution of wealth differs across topologies, and also if the wealth of districts varies across topologies. To do this, however, some sort of index reflecting wealth distribution is necessary.

Such an index does exist, and is called the Gini Coefficient. It is a measure of inequality that can be applied to both wealth and income. Its values range between 0 and 1; 0 represents perfect equality - where all agents have the same wealth, and 1 perfect inequality - all agents have no wealth except for one agent who has all of the wealth.

In order to understand how the Gini Coefficient works, let us first introduce the Lorenz Curve. An example of the Lorenz Curve can be seen in Figure 3.12. [26] As is evident from the graph, the Lorenz Curve plots the percentage of households (in this case agents) against the percentage of income (in this case wealth[2]). In the context of the simulation, it says for example, 10% of the wealth is in the hands of 30% of the agents. Perfect equality is shown simply as the line of "y = x". This can be understood to mean that 20% of the wealth is in 20% of the agents hands, or 21% of the wealth is in 21% of the agents hands, or, more simply, everybody has the same wealth. Perfect inequality is shown as the blue line. Since the Lorenz Curve illustrates the cumulative distribution of wealth, the line of perfect inequality stays at 0 until the final, wealthiest agent is cumulatively included. At this point it jumps up, showing that 100% of the wealth is in the hands of a sole agent.

The Lorenz Curve can be computed directly from the data held in the simulation, making it a simple yet telling method of analysis. The Gini Coefficient, however, is slightly more complex. Take the area of the triangle under the line of perfect equality to have an area of 1. The Gini Coefficient represents the proportion of that area that lies between the line of perfect equality and the Lorenz Curve, labelled Gini Index in Figure 3.13[16].

[2]Although is wealth differentiated from income in that wealth is a measure of assets and income a measure of inflows and outflows, the Lorenz Curve as well as the Gini Coefficient can be used to assess the distribution of both.

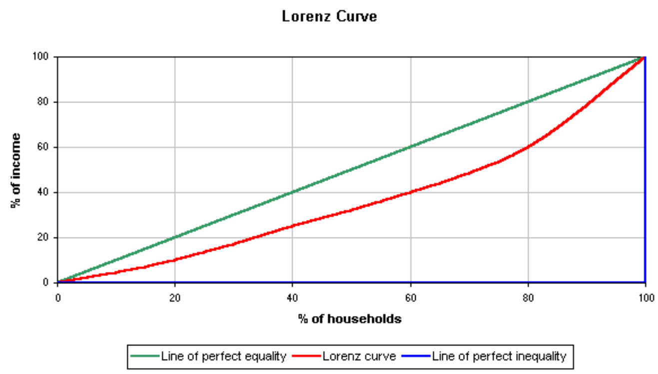

Figure 3.12: An Illustration of the Lorenz Curve

If the area under the Lorenz Curve is B, and the area above (but below the 45 Degree line) is A, then the Gini index corresponds to A/(A + B). Since A + B = 0.5, G = 2A = 1 - 2B. Since B is the area under the Lorenz Curve, if its function is known, the Gini Coefficient can be found by integration.

$$G = 1 - 2\int_0^1 L(X)dX$$

However, in the context of this simulation, it is infeasible to determine the function of the Lorenz Curve, so instead a method of calculating the Gini Coefficient directly from data will be used.

For a random sample S, with values y_i, $i = 1\ to\ n$, indexed in increasing order ($y_i \leq y_{i+1}$) we can compute G(S), a consistent estimator of the Gini Coefficient to be:

$$G(S) = 1 - \frac{2}{n-1}\left(n + 1 - 2\left(\frac{\sum_{i=1}^{n}(n+1-i)y_i}{\sum_{i=1}^{n} y_i}\right)\right)$$

By a consistent estimator, we simply mean one which converges in probability to the true value of the parameter as the sample size is increased.

Having established a way of assessing the distribution of wealth, let us move on to see the distribution found in the initial model. Across the different network topologies, there was only a small difference in the distribution of wealth - measured as the Gini Coefficient. This is illustrated in Table 3.3. This small difference across varying topologies called for a test into its significance in order to determine the probability of a measurement occurring by chance. Since it is necessary to compare 4 samples, it is not appropriate to use a test statistic such as the t-test. Not only would it increase the amount of computation, but in addition, the type one error rate rises with the number of tests we perform. A type one error refers to a false rejection of a true null hypothesis. Instead I will use the ANOVA test, a generalisation of the t-test to cover more than 2

Figure 3.13: Diagram indicating the Gini coefficent

groups, with a hypothesis H_0 that the means do not differ, and H_a that they do.

In order to perform the test, the following steps are taken:[32]

1. Calculate the sample average for each group
2. Calculate the average of these averages, $\bar{x}$
3. Calculate the sample variance of the averages, S^{*2}
4. Calculate the sample variance of each group
5. Calculate the average of all sample variances, S^2
6. Calculate the F Statistic:

$$F = \frac{nS^{*2}}{S^2}$$

   where n is the number of items in a group.

The F-Statistic is 0.3, which is well below the necessary threshold, meaning that it is not possible to reject the null hypothesis that the means are equal. This in turn means that the topology does not have a statistically significant effect on the *global* distribution of wealth. Perhaps it is again a question of the lack of necessity and benefit in trade. Since most agents are self-reliant, the problem of redundancy of network topology is presented. In essence, it is not possible to make proper comparisons based on topology due to the lack of influence it has on the evolution of the simulation, since production does not require a network.

However, in spite of this, it is possible to see the effect of network topology in closing the gap of wealth across districts. Figure 3.14 shows the wealth per

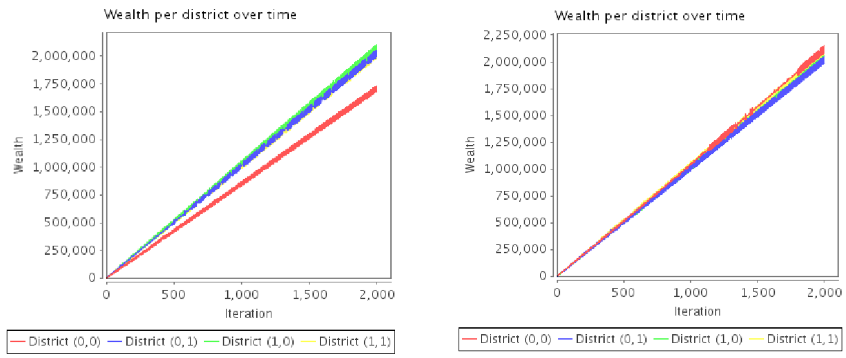

(a) Wealth in a Local Disconnected Network (b) Wealth in a Small World Network

Figure 3.14: Wealth With Time

Topology	Gini
Local Disconnected	0.183
Local Connected	0.195
Small World	0.192
Global	0.186

Table 3.3: Average Gini Coefficient across topologies

district of a Local Disconnected Network, and of a Small World Network. It is clear to see that the wealth in a district is affected by topology when the prices used to calculate wealth are average prices *per district* rather than global ones. Although it would be possible to compute the Gini Coefficient on a per district basis, the lower number of agents involved in the calculation can result in the Gini Coefficient becoming extremely skewed, hence graphical analysis was used instead. [3] The differences seen are indicative of the fact that allowing goods to propagate globally through the network leads to more even prices and hence more even value of goods, which in turn leads to less differentiation between districts based on wealth.

Having established that the Gini Coefficient - the measure of inequality - is not significantly different across topologies, it is necessary now to ask whether the value correlates to the real world. In short, it does not. This self-reliant world is a picture of harmonious equality. It does not correlate to values we witness in life. In reality, Gini Coefficients rarely fall below 0.25 for income distribution, which is almost always lower than wealth distribution. Countries such as the UK generally fluctuate between 0.3 and 0.4, and the USA 0.3 to 0.5. In reality, wealth distribution tends to follow a *Pareto distribution*. This

[3] Although the Gini Coefficient could be computed for each district by having more agents per district, my focus lies more in the global distribution of wealth than the local, since this global focus is more closely related to development in economies which will be discussed later.

is sometimes known as the *Pareto principle* or the *80-20 rule* which says that 80% of the wealth is controlled by 20% of the population. [4]

From simulations using this model, the Gini Coefficient rarely reaches above 0.2. Although the initial endowments and production functions have a uniform distribution, it would still be hoped that the Gini Coefficient would have evolved to become more realistic. Simulations of 20000 iterations were conducted to check that the simulations were being run for a sufficient length of time, but little changed and in some cases the Gini Coefficient even decreased. This must be investigated further prior to being able to suggest why this might be the case, and will indeed be a topic in Section BLAH.

3.3.4 Conclusion

The initial model paints a world in which production is largely optimal. An agents fortune is determined simply by his proficiency in the production of goods. The graph in Figure 3.15 illustrates the small amount of trade apparent in the simulation, with the largest average percentage of trade being just 3.4%. However, it is also clear that as we move from the autarky to the Global Network, there is a clear trend of increasing trade. This again highlights the importance of "open borders" in allowing the realisation of trading opportunities. However, despite this correlation, the amount of trade is still extremely low, so low that it is difficult to evaluate much regarding networks and the impact of crossover agents.

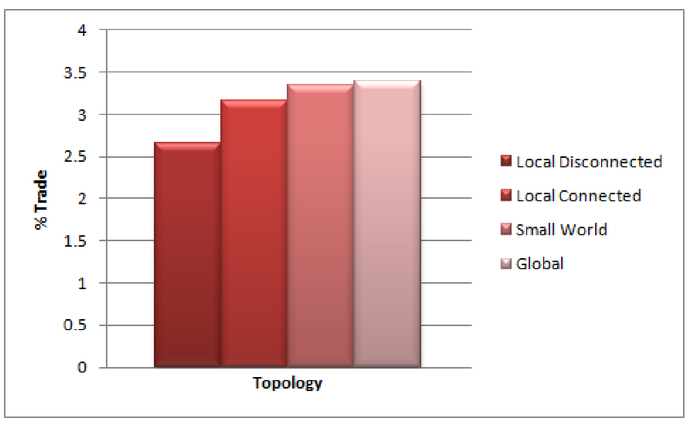

Figure 3.15: The percentage of trade across different network topologies, averaged over 10 simulations of each sort.

The model offers insight into the reasons for emergence of specialisation. It allows agents to be characterised by the strategy they adopt and offers sound conclusions as to what causes certain strategies. The specialisation in the model is not unlike the way the world works. People (or companies, and even countries)

[4]The Lorenz Curve is actually a function of the CDF of the Pareto distribution.

produce what they are most proficient in producing. This is not only reflective of a capitalist economy, but is also an argument for increased efficiency with globalisation. By widening the scope of people with whom agents interact, it is possible to achieve opportunities in efficiency that would not be possible in isolated economies (such as the Local Disconnected Network). By widening this scope, the poorest agents, who are generally pure traders, can nonetheless increase their wealth by a factor of 2. Open borders or open trade allows the specialisation of agents - it allows pure producers to rely on demand for their goods and traders to rely on supply by a larger number of the population. It is the ability to specialise, as Adam Smith noted in his discussion on the division of labour, that contributes to greater efficiency.

However, the distribution of wealth realised by all networks was incomparable to anywhere in today's world. This conflict with what is witnessed in reality is worthy of investigation. In addition, it is difficult to judge how the network topology affects trade and the spread of goods when the vast majority of agents produce, hence the networks are almost redundant. Since the importance of topology in the study of complex systems cannot go unnoticed, trade must be increased in order to delve deeper into the effect of topology on efficiency, globalisation, allocation of resources and hence wealth distribution. One of the major downfalls of economics is the tendency to be concerned more with the actions of individuals than with the interactions *between* them. This is contrary to the view of complex systems research, where it is considered that a key point of interest is the interactions between elements of a system, and the structure of the network of interactions as opposed to the elements themselves. In this context the cooperative interaction (so far) is solely trade and the structure is the topology imposed on the system. To gain insight into the robustness and efficiency of these networks together with their emergence, trade must be increased. Not only is this necessary for proper evaluation, but it is also of fundamental importance if we are to progress in a direction that better reflects macro trends that are witnessed. It is not the case that the vast majority of us are self-sufficient; we do not farm our foods and do not make our own computers. In reality, produce comes from a few wealthy, able firms who are proficient in the production of a certain good. Another thought is that wealth increases linearly and so does the stock of goods. Perhaps this influences the convergence of prices over time - in reality efficiency decreases cost and hence (for many industries) prices for consumers, and excess demand or restricted supply cause price rises. Perhaps consumption of goods could influence the trend of prices over time.

Chapter 4

Increasing Trade

4.1 The Method

The world of production stemmed from the fact that the sacrifice involved in trade decreased its benefit relative to surrender free production. However, we also know that the symmetry of the Cobb-Douglas utility function means that the goal of maximizing utility is best achieved by not only increasing the stock of goods, but also balancing out the stock of goods owned. I needed to find a simple way of tipping the scales to favour trade more frequently, whilst still retaining or even increasing the similarity to the world we live in.

The answer is to restrict the production of both goods. To do this, I permitted a configuration option in which you can specify the percentage of agents who are able to produce both goods. For this subset of the population, half would be able to produce only Good 1, and the others only Good 2. The result of doing so is really quite astounding. It is important to note that an agent may be able to produce neither of the goods. This was allowed since I believed it would be interesting to introduce some true merchants that must make a living solely through price arbitrage.

An important question is whether or not this better reflects reality. I believe yes. It is easy to see that it is not the case that everybody is able to produce everything and some people can't produce anything. In industry for instance, companies specialize. A trainer making company does not manufacture bottled water on the side. Although the simulation does not model firms, it is fair to say that not all agents should have expertise in the production of both goods. By restricting the number of agents who do have expertise in both goods, we create an inherent demand for these goods - trade is necessary for some agents. So not only will it paint a better picture of reality, but it will also allow us to see how the various topologies cope with this new load. In addition it will provide an opportunity to see if there are benefits to being a crossover agent in this new world.

In this chapter I investigate the effect this alteration has on the evolution of the simulation, including the effect on wealth distribution, specialisation and price dispersion. In addition, the trends exhibited in varying the percentage of agents who are able to produce both of the goods will be studied and evaluated.

4.2 Evaluation

The percentage of agents who were able to produce both goods was incremented in 10% intervals, from 0% to 100%. For each interval, 10 simulations were run of the local disconnected, local connected, and small world networks. The global network is excluded herein since it's results are extremely close to that of the small world and local connected network, however, it's search costs are massively increased. This unrealistic and extremely costly topology is therefore of little use in relating to the real world, and due to time constraints, no further investigation into this structure will be carried out.

The following section is divided into 2 areas. One evaluates the effect of increasing trade on macro trends such as prices and wealth, the other takes a more detailed look at the specialization of agents that could lead to these phenomenon.

The aim of restricting who could produce both goods was to facilitate an increase in trade - to make trade more appealing than production. Perhaps then we should first assess whether or not this worked. In fact, from 100% to 0% trade was increased 10 fold - it was a success. In addition, as the graph in Figure 4.1 shows, this increase wasn't linear as may have been expected.

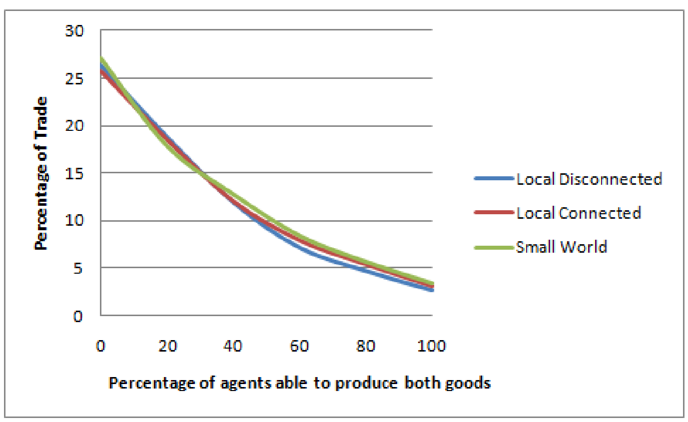

Figure 4.1: Decrease in trade as the percentage of agents able to produce both goods decreases

To better understand this trend, it is necessary to clarify precisely why this restriction lead to more trade in the first place. Recall that the symmetry of the Cobb-Douglas utility function means that balanced stocks of goods maximizes utility for agents. Therefore agents will try and keep their piles of goods as equal as possible. This in turn means that even if they are proficient at producing one of the goods - say Good 1, eventually acquiring Good 2, even if it means sacrificing some of your Good 1 will lead to a higher gain in utility as it will lead to more balanced stocks. Therefore, if an agent can't produce Good 2, eventually it will need to trade for it.

However, the non-linearity is indicative of the fact that the more agents that are in a “worse” position, the more trading opportunities exist for these agents. However, after careful analysis, it was clear that this increase in trade was not happening through all agents simply beginning to trade, and this is discussed in the following section on specialization.

4.2.1 Inspecting Agent Behaviour

Specialization

The breakdown of how agents specialize is probably one of the most interesting results from increasing trade. Recall the four categories, pure producers and pure traders, agents who produce or trade respectively more than 99% of the time; and heavy producers and traders, agents who produce and trade more than 50% but less than 99% of the time. You would expect that pure and heavy producers move to the pure and heavy trader categories as the percentage of trade increases. However, this is actually not the case.

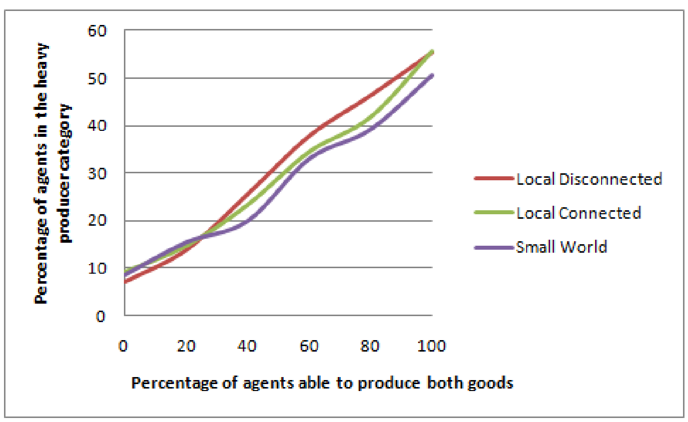

Figure 4.2: Change in percentage of heavy producers as the percentage of agents able to produce both goods decreases

Recall that pure producers and traders were the most specialized with respect to the action they performed most frequently. In fact, when the amount of traders able to produce both goods is low, agents move to the two ends of the continuum of specialization. The percentage of agents that fall in to the heavy producer category decreases consistently for all network topologies as the percentage of agents able to produce both goods moves to 0%. This is shown in Figure 4.2.

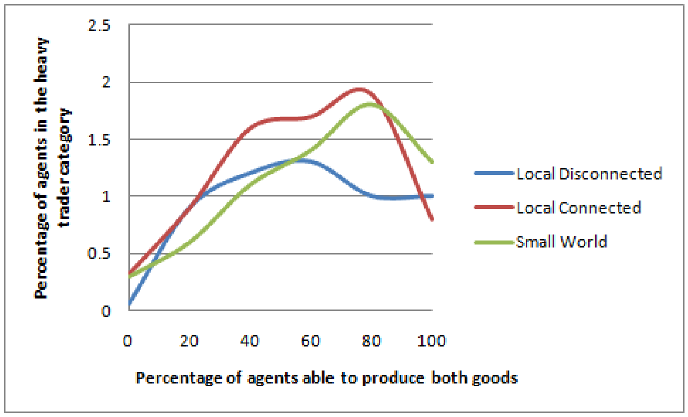

Figure 4.3: Change in percentage of heavy traders as the percentage of agents able to produce both goods decreases

As for heavy traders, all topologies (although at different times and to varying degrees) experience an increase in the percentage of heavy traders to a maximum point before beginning to decrease consistently up to 100% as illustrated in Figure 4.3. Also note that this category seems to show the larges variance between topologies, however it is important to note that the range of percentages is around 2% thus the differences are exaggerated on the scale.

Interestingly, the percentage of pure producers decreases as the percentage of agents able to produce both goods increases up until approximately 80% when it begins to decline (although at a much slower rate). The percentage of pure traders decreases sharply, before increasing two fold at 22%, and then decreases again at just as sharp a rate until the proportion of agents unable to produce both goods reaches 100%. The graphs for pure producers and traders can be seen in Figure 4.4 and 4.5 respectively. Notice also that the pure trader category is subject to the least variance between topologies.

Now the trends have been established, a few more tricky questions to deal with have emerged, and the following seeks to explore candidate solutions to these questions. The questions to be addressed are:

- Why is the number of pure producers increasing?
- What prompts agents to become more specialized?

The fact that the number of pure producers falls as more agents can produce both goods is interesting since the overall amount of trade is decreasing - and pure producers often never trade. At the 100% threshold we characterized pure producers as falling in to one of two categories. Those who alternated between the production of the two goods in a one period cycle, the self sufficient sector. There were also those who were proficient in the production of only one of the goods. These virtually always produced the good they were proficient in, and as such, would often make great trade partners. These agents relied on the need of other agents to trade with them. They did not sacrifice their time to trade, they didn't need to - someone was always willing to trade with them. This fact actually permits answering the first two questions in tandem.

As the amount of agents able to produce both of the two goods decreases, several divides emerge. Firstly we have the divide between agents who can produce both goods, one of the goods, and neither. Obviously those who can produce neither (few and far between) have no choice but to become "shipppers". They make a living by buying cheap in one round, and selling dear in the next. In effect they play the market. However, why do we have less heavy producers and traders and more pure producers and traders? If you recall, heavy producers and traders were characterized by producing or trading for a few rounds and then executing a trade, or producing for one round respectively. Since there were very few heavy traders in the first place, and virtually 0 when no agents were able to produce both goods, the changes in this category have little effect on the overall trend, hence the focus instead lies with the heavy producers. So it is necessary to explain only why being a heavy producer becomes so rare.

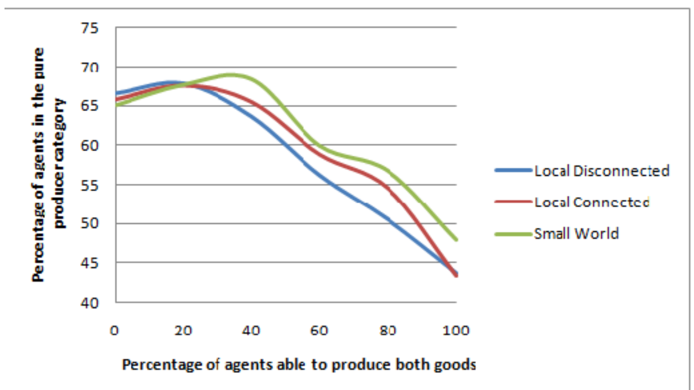

Figure 4.4: Change in percentage of pure producers as the percentage of agents unable to produce both goods increases

It is simpler to understand if you imagine the case when a large percentage of agents are unable to produce both goods. Agents now only fall into a few main categories. As mentioned we have "shippers". In addition, we have agents who have a low production function for one of the goods. These agents generally become pure traders. Trade is almost always more advantageous due to their poor ability in production. We also have agents who are able to produce one of the goods with medium to high proficiency. Even though these agents can only produce one of the goods, the demand for goods in trade is far higher, so they are in a better position to depend on agents initiating trade with them. This happens more frequently in the small world network than in the local disconnected network which is illustrated in the graph by the small world network line being above the other two. This is most likely due to the fact that trading opportunities exist across groups, hence pure producers can rely on demand for their produce from a larger subset of the population, be it by direct or successive trades. However, the number of heavy producers falls as some move to pure production due to the new demand for goods that are only acquirable through trading, and some move to become pure traders as they can, more often than not, find a suitable trade partner. Pure producers don't need as high production proficiency in order for agents to initiate trade with them, and pure traders don't need as low proficiency in order to make trade more beneficial since more agents have skewed stock piles so the probability of finding a trade partner is increased. In essence, the profile or characterization of agents in the extreme specialization categories changes. By forcing agents to specialize in what they produce, the dynamics of the simulation change, and this results in changes in specialization between interactions performed (production or trade).

This said, there are a few interesting points on the graphs worth noting. Firstly, as briefly mentioned, the percentage of agents who are pure producers decrease between the 20% of agents being able to produce both goods and 100%, as shown in Figure 4.4. Before this point, for all topologies, the percentage of pure producers begins to increase gently. In addition, the orders are reversed. Before, the disconnected network had the most pure producers and the small world network the least. This decline is indicative of some form of saturation. The peak proportion of pure producers has been reached, but this does not answer why it begins to increase. One possible explanation is that the number of agents relying on agents initiating trade with them can't pass this peak at 20% as otherwise there aren't enough agents in the rest of the population to support their need for "free exchange". Thus the agents who aren't as proficient producers in one of the goods can no longer survive as a pure producer as the percentage of agents able to produce both goods decreases from 20% to 0%.

A second interesting point relates to the graph of the pure traders, shown in Figure 4.5. Here, we have that at first the proportion of pure traders decreases sharply, and at 80% doubles with just as steep a gradient, before decreasing again afterwards.

In addition, still looking at the pure traders graph in Figure 4.5, we see that the effect of topology on the percentage of agents in this category is negligible. This is most likely due to the fact that trading is still a last resort for agents. Thus, the topology may effect the wealth of the pure traders, but it has little impact relative to production functions on who will be a pure trader.

In conclusion, as the percentage of agents able to produce both goods is decreased, the profiles of agents in each categorization differs, and this is partly due to how the utility function influences the optimal strategy of agents, but in addition, due to the new dynamics that emerge as a result of this forced specialization.

This movement to the edges of the continuum is in fact a better reflection of reality. We actually have very few industries in which the level of production and trade is even. Many companies produce goods and make sales, and others buy goods purely to sell on, such as cash and carries. There are few who both produce their own goods and buy other goods to sell on. The best example of

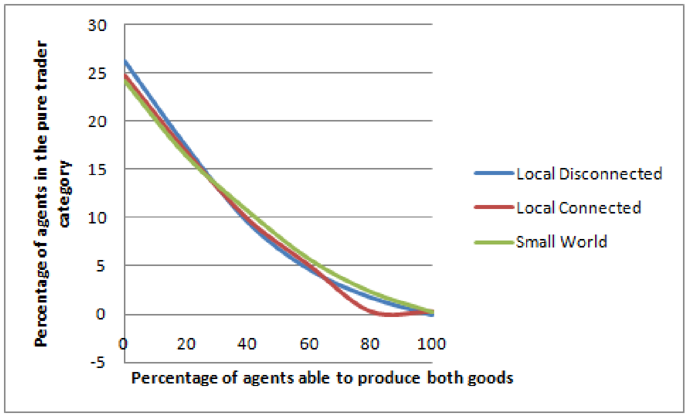

Figure 4.5: Change in percentage of pure traders as the percentage of agents unable to produce both goods increases

these companies is most likely supermarkets who buy branded goods as well as having their own brand. However, on the whole, and even in the business models of supermarkets, there is a clear distinction between merchants and producers. Thus this trend in the model is quite realistic for the simplicity in the economy being modelled.

4.2.2 Global Trends

Price Dispersion

The increase of trade had no consistent effect on the prices of goods across populations. However, the effect it had on the dispersion of prices was dramatic. The graph in Figure 4.7 shows this trend for each of the network topologies. For the small world network the standard deviation of prices was relatively consistent across the different thresholds. The amount of trade still leads to solid price convergence across all districts, with consistently low deviation. This highlights the efficiency in goods spreading globally through the network even when the amount of trade is high, and the stock piles of goods more inclined to become uneven. Since the path length is short, goods can reach every agent quickly and hence there are few trades between groups that go wanting, thus the range of prices is lower. In contrast, in the local disconnected network an increased need for trade and thus an increase in trade, leads to many opportunities between groups that are not realized. This forces the prices between groups further apart. This is shown in Figure 4.6

For instance, in one district perhaps there are many agents who can not produce Good 1. Therefore the price is excessively high, since the demand for Good 1 is far higher. This district is unable to trade with another district where the price of Good 1 is low, and therefore the gap exists. This could explain why the increase in trade leads to sharp increases or decreases in prices

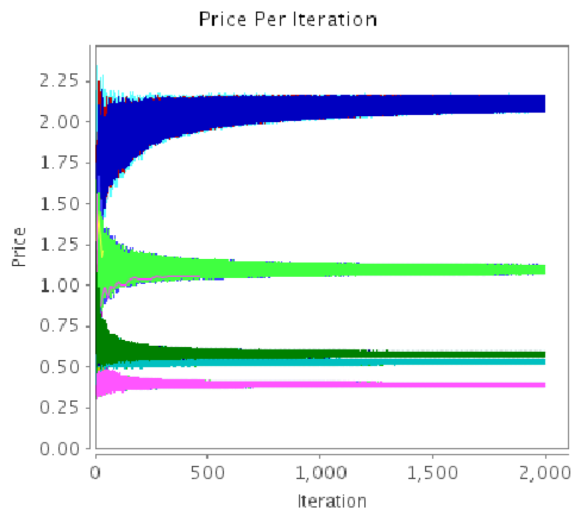

Figure 4.6: Illustration of huge price dispersion at the 100% threshold for the local disconnected network

across districts. Although this accounts in the increase in price dispersion, an interesting point on the graph in Figure 4.7 is the point at which standard deviation in the local disconnected network plummets from 0.37 to 0.21. This occurs when the change in the percentage of agents able to produce both goods increases from 40% to 60%. Past 60%, it plateaus, before falling again. It is as if 40% is a threshhold, after which the combination of the demand for goods only acquirable through trade, and the autarky imposed by the local disconnected network result in a massive increase in standard deviation, or price dispersion. This is the point at which the scales tip, meaning the limitations of no cross district trade are exhibited in the form of some districts suffering massive prices, and others extremely low prices.

Wealth Distribution and Global Wealth

As may be expected, there is a clear correlation between the percentage of agents unable to produce both goods, and the Gini Coefficient - a measure of wealth inequality introduced in the previous chapter. This relationship can be seen in Figure 4.8. As the percentage of agents able to produce both goods increases, the distribution of wealth across all topologies becomes more even indicated by the decrease in the Gini coefficient. The variation between topologies is still virtually negligible.

In addition, the higher in wealth inequality actually reflects a far more realistic world. In reality, the Gini coefficient for areas with respect to income distribution such as the UK and the USA is approximately 0.35. Thus it is

Figure 4.7: The effect of increasing the percentage of agents unable to produce both goods on standard deviation

fairly realistic that in a free market, not so dissimilar to that of the UK, this level of wealth inequality is realized. However as previously mentioned, wealth inequality is thought to be quite a lot higher than income inequality so perhaps there is still room for increasing the coefficient.

So what is causing society to become so much more unequal when the percentage of agents able to produce both goods is low? It can't simply be the unevenness imposed by not all agents being able to produce both goods, since when no agent can produce both goods, the level of inequality is higher still. Over time, as may be expected, global wealth falls as production capacity falls. However, what is interesting is that the wealthiest agents remain relatively unaffected, and the poor just get poorer (when the percentage of agents able to produce both goods is low), thus widening the gap and increasing the Gini coefficient. However, the fact that the wealth of the wealthiest agents remains unaffected, whilst global wealth actually falls implies that the rich are richer relative to the other agents. This further reinforces the justification as to why the distribution of wealth is more uneven when less agents can produce both goods. Even though this explains the uneven wealth distribution, why the poor become relatively poorer, and the rich relatively richer is still unexplained.

The rich in society are still pure producers. They also have an increasing number of agents wishing to intiate trade with them. Thus they get more free exchanges when the percentage of agents able to produce both goods is low. As a result, the rich are becoming richer as their product is in demand. This is analogous to the theory in economics known as *wealth condensation*. This states that there is a correlation between being rich and earning more - new wealth condenses to the already wealthy. In the context of the simulation, there is a correlation between being proficient in the production of a good, and becoming wealthier as a result. This is trivial to see, since if you can produce more, you will be wealthier. However, what is less trivial, is how increasing trade affects

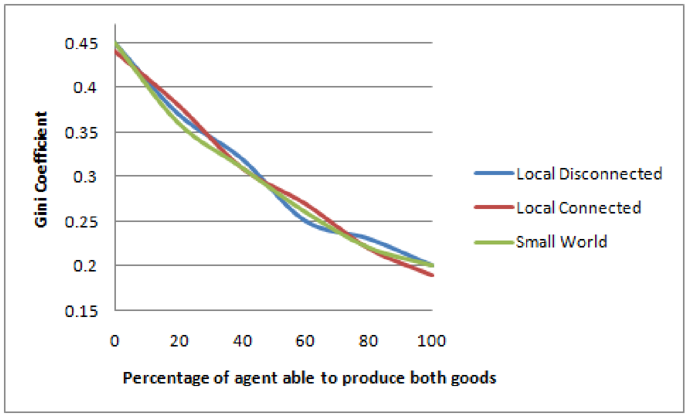

Figure 4.8: The effect of increasing the percentage of agents unable to produce both goods on wealth inequality

the distribution of wealth. Now, if you are proficient in production of a good, you will become even wealthier than when there is less trade. This is due, as mentioned, to the extra demand for goods since less agents are able to produce both, more are forced to trade for it. This could be seen as unfair allocation of resources, however, on the other hand, perhaps it is reasonable to expect that in a free market, those who can do something well reap the rewards of their ability. Perhaps it would be more unfair if they did not become wealthier for providing a good service.

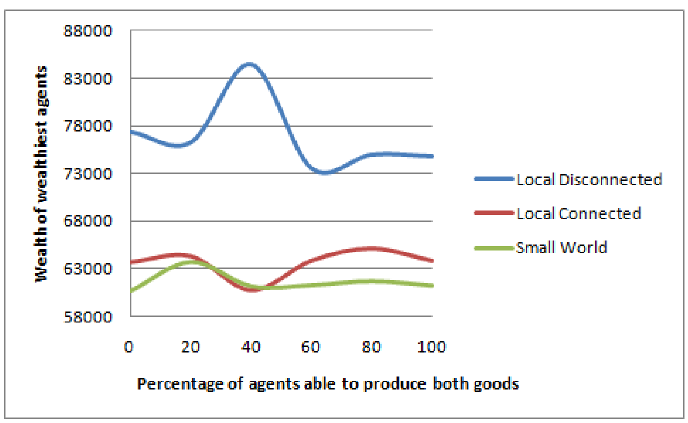

Figure 4.9: The effect of increasing the percentage of agents unable to produce both goods on the wealth of the wealthiest agents

Now let us examine why the poor get poorer when less agents produce both goods. Recall that for pure traders, trading was a last resort. Despite the increase in trade, being a pure trader is still a last resort. Those who can rely on agents initiating trade with them become pure producers and enjoy the wealthier life. Traders however, are those who rely on the pure producers for trade. As discussed, new wealth condenses to the already wealthy, thus, this wealth is not available for the pure traders. So how does wealth *condense* to the already wealthy? Since pure traders are constantly surrendering goods in trade to their trade partner - the pure producers, they are constantly increasing the wealth of producers and simultaneously hindering the growth of their own wealth. They simply survive.

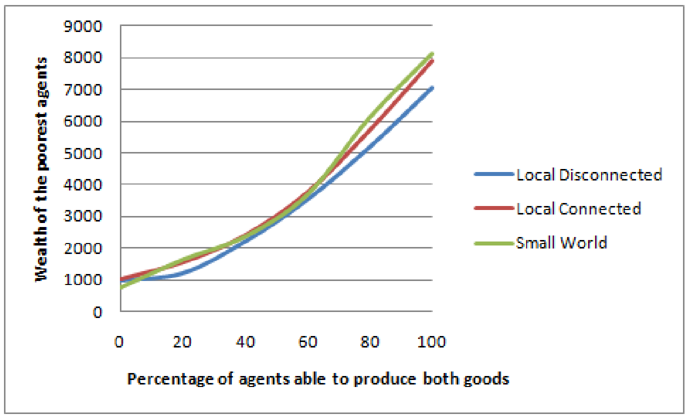

Figure 4.10: The effect of increasing the percentage of agents able to produce both goods on the wealth of the poorest agents

The graphs in Figures 4.9 and Figure 4.10 show wealth of the wealthiest and poorest agents respectively against the percentage of agents able to produce both goods. An interesting point on the graph depicting the wealth of the wealthiest agents in the sharp decrease in wealth when the percentage of agents able to produce both goods increases between 40% and 60%, for the local disconnected network. Incidentally, this is the same range where the massive increase in standard deviation occurs, so it seems that this is an important tipping point in the local disconnected network.

So as the amount of trade increases, more dependencies emerge, and the emergence of wealth condensation favours the rich and punishes the poor in an increasingly divided world.

4.2.3 Conclusion

In conclusion, the increase in trade has highlighted how initial conditions affect the evolution of the simulation. It has shown a progressive movement, redefining profiles of the categorized agents as product specialization is imposed by the system. We also witnessed a tendency for agents to become extreme specializers, and even saw the percentage of pure producers increasing. The effect of a local disconnected network on the dispersion of prices was further exaggerated with the increase in trade, and we saw the sensitivity of the local disconnected network to the increase in percentage of agents unable to produce both goods. The Gini coefficient better reflects reality, and indicates the divide in rich and poor. However, people often suggest that wealth distribution follows the *Pareto Principle* or *80-20 rule*, thus although it is more realistic, it is difficult to judge the difference between income and wealth distributions, even when dealing with countries such as the UK.

Chapter 5

Introducing Consumption For Survival

Consumption is the act of consuming one of the goods. In the simulation, Good 2 is representative of money, which is not consumed. Good 1 encapsulates every other conceivable good, which can be consumed. In the model, agents must now consume a fixed quantity of goods, and if they are particularly wealthy, they consume more. Being unable to consume the required, or baseline, quantity reflects an agents inadequacy in the economy and forces an economic death, which is to say, the agent exits the economy and hence the simulation. This chapter focuses on the questions of why to consume, how it was achieved and what the consequences were. It also looks at the threshold values for the baseline amount to consume - what is a sustainable quantity?

5.1 Motivation

Consumption is inherent in society. It also ties in with the notions of wealth and distribution. In essence, wealth can be a better indicator than income (strictly in this context [1]) when it comes to testing how an agent in the economy can perform when times are bad - does the agent have a stock of goods cushioning it? Consumption in the context of the economic simulation is not akin to the consumption of food and so on, it is more reflective of expenditures to sustain their "business" of production and trade. Importantly this is not necessarily a monetary cost, but rather it is the resources needed - machines, ingredients and so on - so costs are not reflective of money - Good 2 - but it may be necessary to spend Good 2 in order to acquire resources - Good 1.

[1] Its suitability in this context stems from the fact that there are no liabilities - no debt. In the real world, measures of household wealth often class people in the west as poor since they have more debts than assets! Thus, in reality either a notion of income is used if you are studying it on a household level, or a combination, or wealth if you are investigating how good a country's finance sector is etc.

There are several pertinent questions to be addressed at this point. Firstly, why do the wealthier consume more? A poor performing agent in the economy, who is able to produce little of either good, for instance, is analogous to a small business. Its size is reflective of costs, and a small size has lower total costs. Similarly, someone who is particularly wealthy and is extremely proficient in production of some goods is a bigger player in the economy, and thus has higher costs. For example, the cost of running the factories of a massive food producing company supplying to supermarkets are far greater than that of a local farmer growing small amounts of produce to sell This explanation deals with producers, but what about the traders? The traders act in a similar fashion. Less wealthy traders often trade in smaller quantities, and bigger traders in large quantities. The cost of standing on a market stall selling the produce of the local farmer is far less than that of shipping the immense quantity of ingredients from a huge farm to the massive food producing factory for them to make their products.

We might then consider how one can need Good 1 to produce Good 1? Since Good 1 encompasses all other goods, it is a very abstract notion of a good. In the simulation we cannot distinguish between Good 1 being an apple and Good 1 being a lorry, nor is it necessary to do so. All that need be represented is that production and trade are not free, stocks of goods do not increase infinitely and prices are affected by the level of demand. So we can introduce consumption - the exiting of goods from the economy through use -as a new way of finding the agents who can sustain themselves in this world, and a new sort of demand for goods.

Consumption - or resources needed by businesses - is extremely relevant, especially in today's economic environment. Examples of the importance of sustainability of consumption by businesses, and their dependence on price fluctuations and changing levels of demand are in abundance. The most obvious example for price rises is the effect of the rises of oil price on virtually all businesses. Oil underwent huge price increases in the years of 2003 to 2008 - its price per barrel increased approximately by a factor of 4 [11]. Contributors to these price rises are increased demand: as developing economies develop they tend to consume more resources that relate to wealth (for instance oil for cars) and that drive economic growth. This, coupled with the slowdown of petroleum production and other oil-related products means that the difference between supply and demand is furthered and prices increase. The increases affect all businesses using transport - distribution companies, the shipping companies, air travel companies and even restaurants. This is indicative of both the size and scope of the effect that changes in supply, demand and prices of goods have on companies. In this simulation, consumption introduces the notion of scarce resources and looks at the effect it has on the dynamics of the simulation.

5.2 Implementation

A key necessity was for consumption, in the beginning, to be feasible for the large majority of the population. Therefore, the baseline quantity - the necessary amount to consume was computed at the beginning of the simulation.

In the simulation, the baseline quantity was set to be a proportion of the average stock of goods existing as initial endowments in the world. Mathematically:

$$C = a \times \left(\frac{\sum_1^n (g_1^i + g_2^i)}{2n} \right)$$

where

C is baseline consumption, a some constant less than 1, n is the total number of agents, g_1^i, g_2^i are the stock of Good 1 and Good 2 that agent i is endowed with respectively.

This baseline was the amount of Good 1 agents had to consume on every *mathematics* iteration. It was necessary to vary a and n in order to find stable values for both of them, so that some agents were still affluent, some agents were poorer, and agents could (for the most part) still exist in the economy for prolonged periods while simulataneously having consumption influence the simulation in some way.

The extra amount an agent is to consume is set to be a proportion of their current stock of Good 1:

$$E = b \times (g_1^i + (\frac{g_2^i}{p}))$$

where

E is the extra amount of Good 1 to consume, b is a constant less than 1, g_1^i, g_2^i are the current stock of Good 1 and Good 2 respectively held by agent i and p is the average price of Good 1 at the point in time.

Let us go on in the sections that follow to look at the effect of changing these variables, a and m, on the evolution of the simulation, to find stable points and to use these points to evaluate the new world with consumption.

5.3 Evaluation

5.3.1 Evaluating values for constants

In order to implement consumption and ensure a balance between a stable simulation and being able to see the effect it had on the evolution of the simulation it was necessary to perform experiments to attempt to find stable points. There is in fact a fine line between a chaotic simulation riddled with bankruptcy and a stable simulation with little effect on anything but wealth. The effect on wealth is easily explained by the fact that consumption means goods leaving the system and thus the total stock of Good 1 will definitely be reduced, and stock piles will not grow infinitely thus wealth is considerably lower.

In order to find a stable point, two things were varied: the value of the constant a used in calculating the baseline amount to consume, and the frequency of consumption, m - agents had to consume every m iterations. The value of a was set to be *0.03, 0.06, 0.09* and m was set at *1, 3, 5*. Every permutation of the following was tested. Each simulation ran for 2000 iterations, with 20 districts, 20 agents per district, 2 crossover agents (if applicable) and 20% of the agents able to produce both goods. Simulations were run for the Local Connected, Disconnected and Small World Networks. Results for each were averaged over 10 simulations. The values for the baseline quantity to consume are given in Table 5.1 with the corresponding a value.

a	Baseline
0.03	1
0.06	2
0.09	3

Table 5.1: Table showing value for a and corresponding baseline to consume

Four areas were examined in order to evaluate the effect of the values that m and a had on the evolution of the simulation, namely, the level of trade, average wealth, wealth distribution and price. In addition, significance testing indicated that the differences across topologies were insignificant and thus the topologies will not be examined in isolation.

The results revealed that the largest contributor to the evolution and stability of the simulation was the value chosen for m. The amount of time between agents having to consume consistently resulted in dramatic differences in the ability of agents to cope with the requirement of consumption.

Beginning with m being set to one, agents were forced to consume on every iteration. In this way, agents who are unable to save up Good 1 over a period of time end up suffering massively, and those who suffer the most are the pure traders. They are typically unable to produce much of either goods, and as such are forced to purchase the quantity they are to consume. Due to their poor production proficiency in both Good 1 and Good 2, and the fact that they are required to consume on every iteration, they were the first to go bankrupt.

Surely however, they would be able to purchase the goods from agents who are most proficient in production. Unfortunately, agents who are proficient in the production of Good 1 consume more, and are unrealistically less wealthy as a result. A lot of trading opportunities evaporate for the pure traders. Their stock of Good 1 and Good 2 is low, and pure producers have a larger amount of Good 2 than Good 1 - Good 2 is not exiting the system. This large difference can be seen by inspecting the movement of goods for a pure producer illustrated in Figure 5.1. The stock of Good 1 stays relatively constant, and low, whilst the stock of Good 2 grows. Thus, since price between two agents is calculated to be

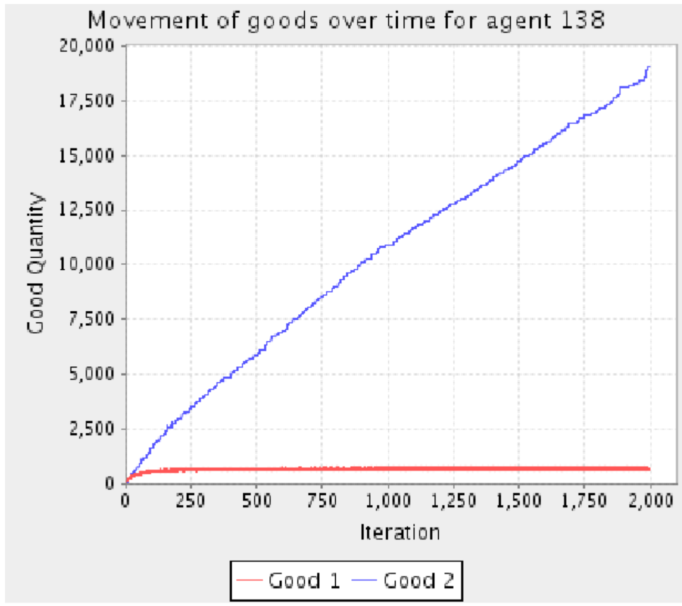

Figure 5.1: Illustration of skewed stock piles, even for a pure producers causing price inflation

$$P_{i,j} = \frac{g_i^2 + g_j^2}{g_i^1 + g_j^1}$$

the price for one unit of Good 1 is extremely high if the quantity held of Good 2 is far greater than the quantity Good 1. To see just how large fluctuations in price are, and to see in addition how high they are (considering without consumption they approximate to 1), see Figure 5.2. This in turn means that pure traders who can't produce Good 1, and can produce only a small amount of Good 2 are unable to afford just one unit of Good 1 hence are bankrupted. This forces the level of trade down significantly for this value of m, as illustrated in Figure 5.3.

The level of trade falls below 1% for $m = 1$ and $a > 0.06$ as pure traders are bankrupted and trading ceases. Pure traders being bankrupted however is not the only reason for this massive decline in trade. In fact, there is also a change in strategy of agents. Pure traders who can produce the required amount of Good 1 migrate to the pure producer category. This can be seen by comparing the two specialisation graphs, mapping the production ability of an agent with their specialisation shown in Figure 5.9. On the x and y axis we have how much they can produce of Good 1 and 2 respecitively.

Each series corresponds to a specialisation category. With consumption (Figure 5.4a), pure producers emerge when production of Good 1 is lower and Good 2 is 0. They can survive by producing then consuming. Pure traders

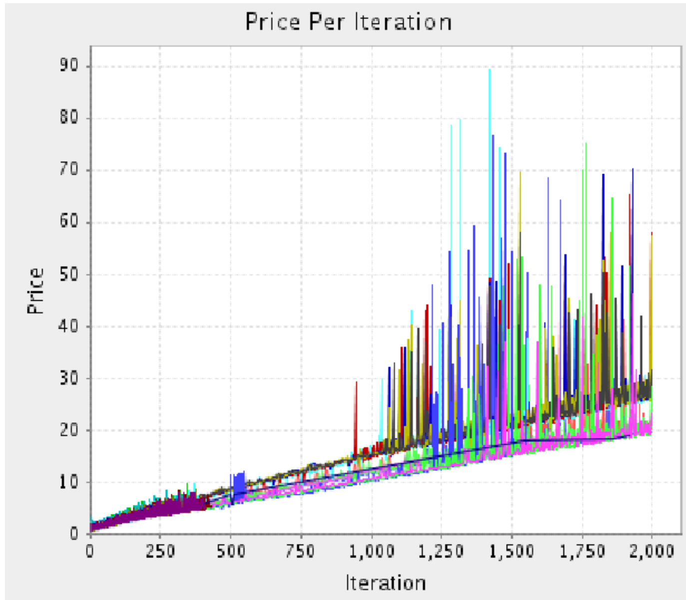

Figure 5.2: Illustration of how skewed stock piles effects prices

are only apparent when they can produce 0 or practically 0 of both Goods. In addition, with consumption heavy producers now occupy the region where Good 2 is produced and Good 1 isn't, or the amount is small. This shows that those agents who used to be pure producers (shown in Figure 5.4b) are now forced to engage in more trade in order to allow them to consume - a strategy shift. Far less agents are in a position to initiate a sale of Good 1 so they can no longer rely on agents initiating exchange with them. These dramatic strategy shifts indicate the new demands that have been inflicted on the popualtion have radically changed the dynamics of micro interactions between agents.

As a result of the miniscule amount of trade, prices seem uneffected by the value of both a and m. However, this is misleading since prices are rarely made thus they seem not to be inflated relative to the other values for m as seen in Figure 5.5. Wealth however, is a very good indicator of how devastating frequent consumption is on these agents. Average wealth falls dramatically to under 1000 units when consumption on every iteration is required. This can be seen in Figure 5.6. By simply increasing m to 3, average wealth is increased by a minimum of 6 times. In addition, when m is at 1, the Gini coefficient is at its lowest - not because society is more equal, rather, the poor exit the economy (Figure 5.7).

The difficulty caused for agents when faced with such frequent consumption also means that the effect of the value of a is negligible. However, when m is set to 3, agents have time to save, and trade is increased ten fold for a low value of a. At this value for m, changing the baseline quantity to consume has a greater

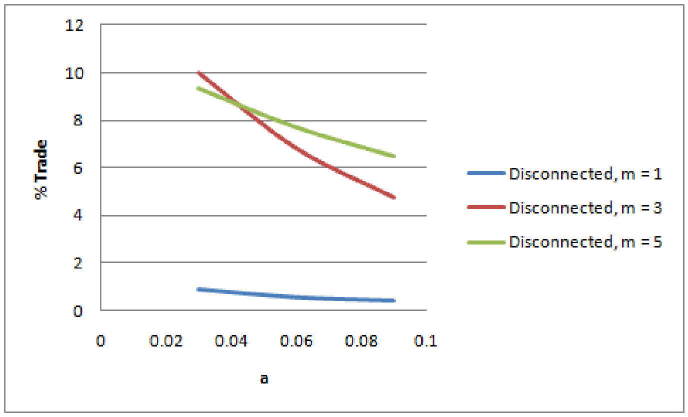

Figure 5.3: Level of trade as m and a are varied

effect. Trade falls sharply as a is increased from 0.03 to 0.09. Two things are occurring here. Firstly, some pure traders are unable to consume this amount, as they can neither produce it, nor afford to buy it. By inspecting the graph in Figure 5.5, it is apparent just how difficult it is for agents unable to produce enough Good 1 to purchase it instead. The average price has risen to over 20 times the price relative to without consumption. As Good 1 exits the system, we witness inflation - the demand for Good 1 is increased and in addition the supply is restricted. Secondly, some pure traders can produce enough to sustain their livelihood, and as such turn into pure producers.

In addition, wealth falls substantiall as a is varied when m is at 3 (Figure 5.6). This value for m is the value at which the baseline amount to consume has the greatest effect. Consumption is manageable, but fragile - it depends heavily on how much agents are expected to consume. An interesting point for when agents are to consume every third iteration is the effect that varying a has on price (Figure 5.5). For other values of a, price is unaffected, however, now it is reduced by half between a being 0.03 and 0.09. What causes prices to fall is fairly misleading. As more agents are bankrupted or move to become heavy producers, indicated by the fall sharp fall in trade, demand for purchasing Good 1 is greatly reduced. There is the more Good 1 available to a smaller population and thus it falls. Imagine at a being 0.03, many agents are competing for Good 1. As such, the stock piles of those agents with a good ability in the production of Good 1 is fairly low as so many agents are buying it from them, which pushes up prices. However, as the price starts much higher when a is 0.09, thus a lot of agents can never afford to buy it and exit the economy, removing a large proportion of the demand. This means supply is less scarce, and producers make less sales. As a result, their stock piles of goods become marginally more even - enough for prices to fall over the long term.

Moving m to 5, the effect a has is dampened. Trade still falls, but both less sharply and to a higher value. Wealth is far greater and doesn't change

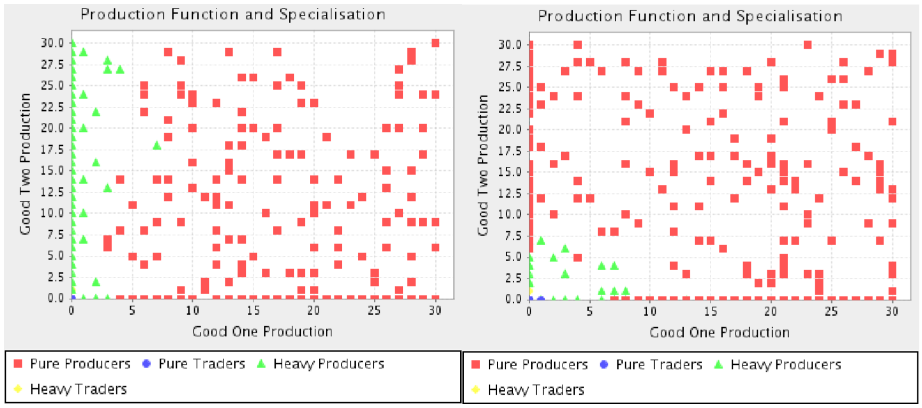

(a) With Consumption (b) Without Consumption

Figure 5.4: Correlation between specialisation and production functions

significantly. It is important to note, that since wealth accounts for price, and the price is lower when m is 5 relative to when m is 3, the increase in wealth when m is 5 is actually of even larger significance since average wealth is higher despite the value of Good 1 being lower.

In additon, the Gini coefficient is higher when m is 5 (Figure 5.7). Despite the difference being small when m is changed from 3 to 5, the negligible variation proved this to be significant via a means test. The reason for the heightened uneven wealth distribution is that their are fewer bakruptcies and so the more poor agents can survive the harsh world. As a is increased, the Gini coefficient decreases slightly as less agents survive.

Another interesting point on wealth is that the growth or creation of wealth in society has a different trend. Without consumption, recall the increase of wealth over time presented a linear trend (although fluctuating with price movements). However, when consumption the frequency and baseline are low enough, this now becomes more logarithmic. This is illustrated in Figure 5.9b. We have wealth creation, however the rate of growth is slower and more realistic. On the other hand, move consumption to be carried out on every iteration and a dramatic change occurs: wealth creation ceases (Figure 5.9a). There is a sharp increase at the same time as a major spike in prices (Figure 5.8), and then wealth remains constant for the entire duration of the simulation. At this point, the new wealth created on every iteration of the simulation is either immediatey consumed or exits as an agent becomes insolvent. This is consistent for all values of a when m is set to 1. This is illustrative of the importance of initial conditions on the evolution of complex dynamic systems.

In conclusion, it is clear that consuming on every iteration when a large proportion of the population are unable to produce both goods is not sustainable and prevents economic growth. It seems the best value is for agents to consume as little as possible as infrequently as often - setting m to 5 and a to 0.03. This

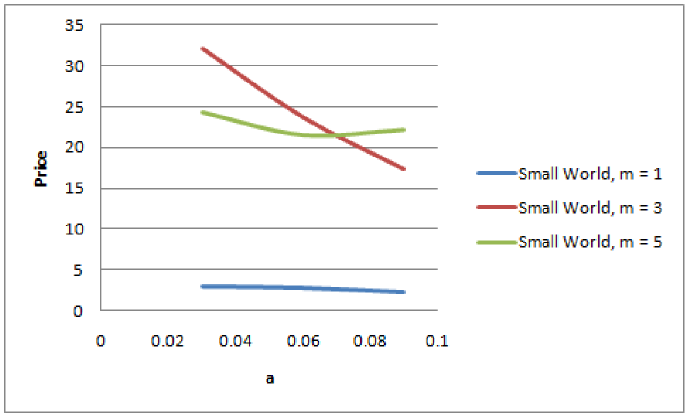

Figure 5.5: Level of trade as m and a are varied

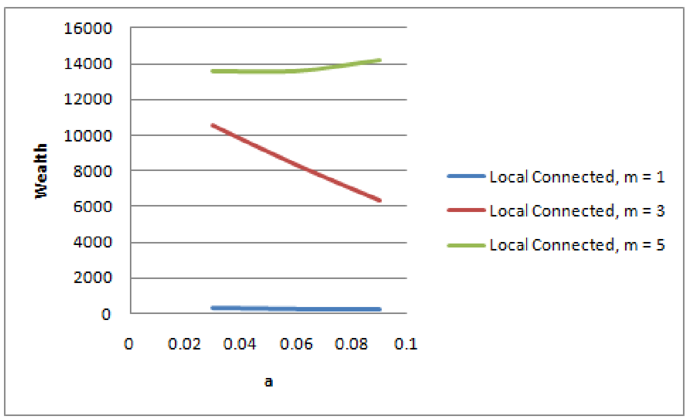

Figure 5.6: Level of trade as m and a are varied

way less deaths occur and the less fortunate agents have a better chance of avoiding bankruptcy. However, one problem is that the wealthy agents perhaps consume too much - leaving supply to dwindle and poorer agents to suffer. In the next section, the additional quantity for the wealthy to consume will be reduced to the same value as a.

With respect to how realistic it is, the answer is, not very! The poor have little opportunity. Perhaps though this is not a downside on actual consumption in the model, but the limitations of not allowing borrowing or financial support for the poor. With banks, pure traders would be able to make massive margins through riding the price fluctuations, however the capital just isn't there for them to begin with. Thus, perhaps introducing an additional interaction of borrowing could allow for traders to engage in trade and exploit the prices - an interesting extension.

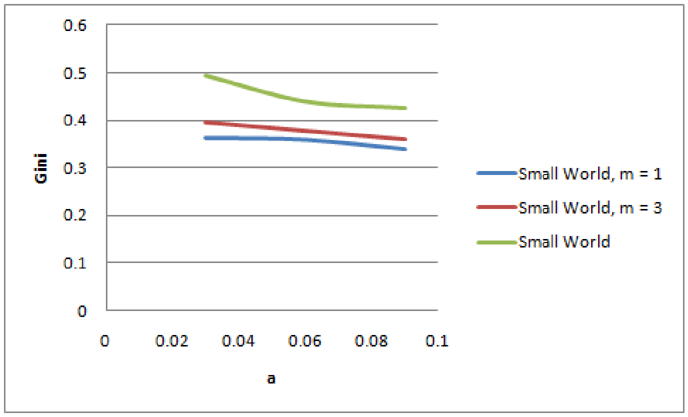

Figure 5.7: Gini Coefficient as m and a are varied

However, the importance of initial conditions was reflected in the changes of straegies witnessed and the price inflation was indicative of the importance of balance between supply and demand. The Gini coefficient was increased as more agents were living on the bare minimum, and less bankruptcies occurred. This can be viewed as analogous to the distribution of wealth of companies: an economy where more firms can survive in a market place irrespective of their size - in other words it is less monopolistic. Although the difference in importance and wealth varies drastically, the fact that small industries are permitted is realistic.

5.3.2 Bankruptcy Chains

Throughout the report an emphasis on the important of networks - the interactions between agents - and their applicability in economics has been suggested but not addressed. In this section, consumption is used to examine the applicability of network theory to the study of bankruptcy chains. A bankruptcy chain in this simulation constitutes the bankruptcy of one agent facilitating the bankruptcy of others which in turn can cause more. It can be thought of as a domino effect. Here, an investigation into predicting whether the dependencies of agents on a bankrupted agent would cause the dependents to also go bankrupt is conducted and compared to the results of the simulation in order to evaluate its usefullness and accuracy. In addition, the applicability of this idea to the field of economics is assessed.

A problem with this investigation is that often, agents are dependent on pure producers, who more often than not survive. Thus bankruptcy is forced upon a single pure pruoducer, and the effect of this on other agents in the population is examined. The pure producer to be bankrupted was chosen based on a heuristic: when consumption is employed the best trade partners are those who can produce alot of Good 1 and little Good 2. Thus an agent fitting this

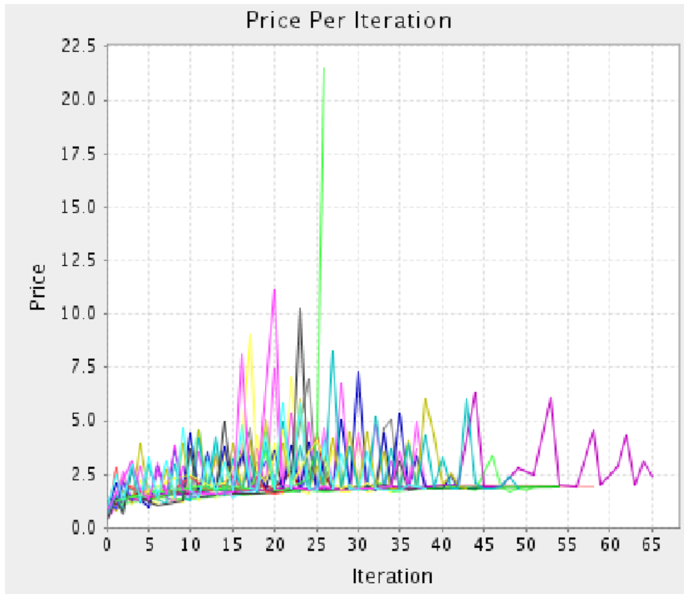

Figure 5.8: Prices for $m = 1$, $a = 0.03$

description is searched for, and the first one encountered bankrupted (removed from the simulation). Deaths and trades are recorded before and after this event to facilitate examination of changes in network structure and resilience for two topologies, the local disconnected and small world networks.

Iterestingly, both realisms and fundamental flaws in the model at hand became apparent upon conducting this investigation. First, the performance of the heuristic will be assessed prior to discussing the results of forced bankruptcy on the local disconnected network and the small world network; finally a discussion on what is severely lacking will follow.

The heuristic is to bankrupt a producer with a high proficiency in producing Good 1 and a poor proficiency in producing Good 2. In order to assess the success of the heuristic, it is necessary to define the aim: to bankrupt an agent who many agents would want to initiate trade with. In order to measure the success of this, a notion of network centrality was employed, specifically *degree centrality*. This measure bases its calculation on the assumption that an important node in network is one which is connected to many other nodes. In the context of the simulation, this is precisely what we are looking for. A matrix of agents in a district is constructed, with every agent labelling a column and row - so for a district of 20 agents a 20 by 20 matrix is constructed. Each position i,j in the matrix represents the percentage of all Good 1 bought by agent i that came directly from agent j. The average percentage is calculated across the whole district and this value is taken to be a threshold representing meaningful purchasing. Each agent in the district is represented as a node in a nework.

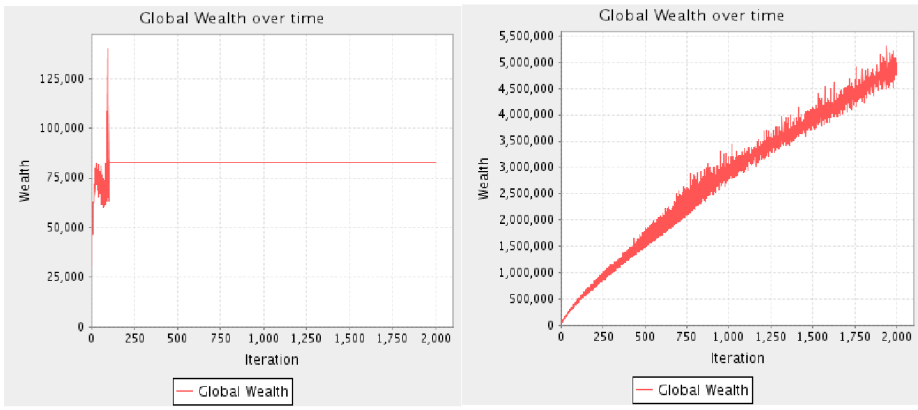

(a) With $m = 1$

(b) With $m = 5$

Figure 5.9: Gloabl Wealth varying m

For each meaningful purchase by agent i from agent j, an arc is added between them travelling from node i to node j. Since this represents purchasing Good 1, it is indicative of a dependence agent i has on agent j to remain solvent. Thus this network can be seen as a network of dependencies. The degree centrality for node j is then calculated to be:

$$C_D(n_j) = \frac{d(n_j)}{N-1}$$

,

where $C_D(n_j)$ is the degree centrality for node j representing agent j, $d(n_j)$ is the number of edges leaving node j and N is the total number of nodes in the network [14].

For all 20 simulations conducted, 80% of the time, this heuristic resulted in the node with the highest degree centrality being subjected to forced bankruptcy and 20% of the time the node with the second highest. Thus the heuristic based method for selecting victimes was successful.

Local Disconnected Network

In the local disconnected network, the fact that each district is isolated from the rest of society means that any implications resulting from forced bakruptcy are self contained within the district that the victim lives. Thus in order to investigate the effect, it is necessary only to examine the victims district.

All the results from the different runs were fairly similar hence for clarity only one will be focussed on. Figure 5.10 illustrates the network produced prior to the forced bankruptcy. The victim agent is agent number 303. As you can see by inspecting this node (red) this has the most inbound edges out of all nodes in the network.

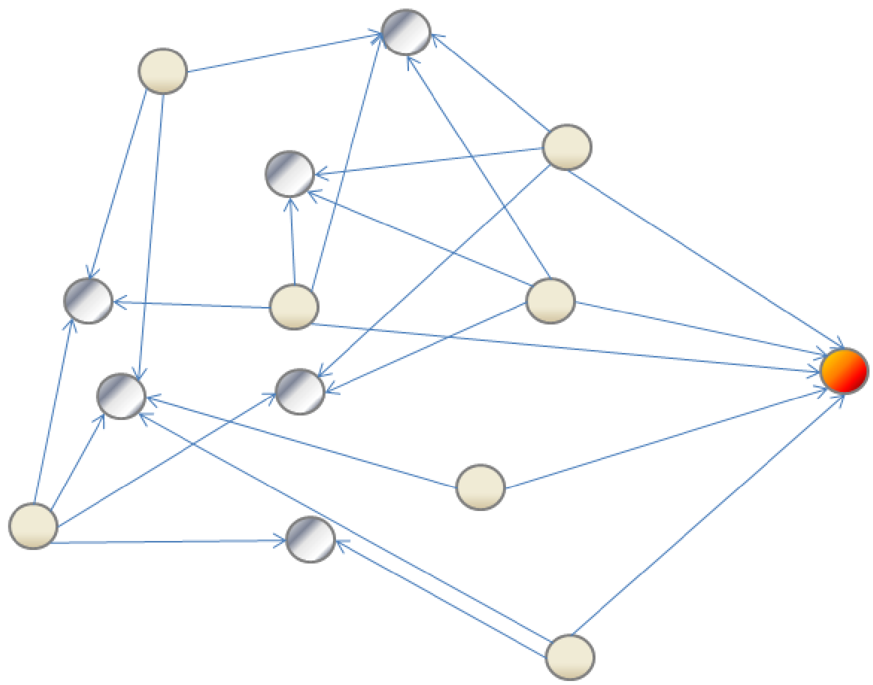

Figure 5.10: Illustration of trade network, victim in red, children in pink, parents in grey

Each child in the network is coloured in pink. These are nodes who only have outbound networks - noone is dependent on them. The grey nodes are dependent on noone. As shown, there are no nodes who have *both* inbound and outbound nodes. This actually presents a large flaw in the model. It is indicative of the fact that traders depend on only producers, producers depend on themselves, and nobody depends upon traders. As a result, upon analysing bankruptcy chains, bankrupting the victim can only ever have an affect on anything directly dependent or transitively dependent on them. Since transitive dependence doesn't exist, it can only ever affect the nodes siblings. This lack of transitive dependencies brings to light an big flaw of the model. It shows that traders are not only poorest and most vulnerable to supply shortages, but also that they, never make enough profit to sell on a good they buy and this prohibts trade chains. They trade only for survival - they do not make large margins and they cannot exploit their position to do so. The simulation does not value traders.

As a result, despite the forced bankruptcy being damaging to the direct buyers, the damage stops their. This is massively unrealistic. Take the current recession to be an example. An initial mistake in the financial sector, specifically "mortgage backed securites" (asset backed securities where cash flow is backed by a collection of mortgage repayments) caused the bust of the housing market. This in turn caused home depot stores to suffer, as well as building industries. Banks suffered huge losses and grew cautious of lending money, meaning investment was hampered through lack of getting a loan. This has a knock on effect in economic growth, unemployment has risen and much more. The point being that despite the simplicity of the model economy being used, you would expect some domino effect. As it is, we have a single bankruptcy effecting only its children simply due to the fact that these children have no children of their own. The path length for goods moving from one agent to the next is too short. This is due to the fact that noone, not even pure traders themselves, wish to buy from other traders. This results in having a network in which there are a few centers of gravity - nodes who have lot of dependents and are important in the network. These attractors in the network are rarely connected to each other, nor are their children connected to other children. There are simply nodes who have children and these children have multiple dependencies.

However, despite the “chains” having a depth of one, the victim’s bankruptcy did have an effect on some of its dependents 100% of the time, and also on others who suffered from the change in the network structure. Hence we shall move on to discuss exactly what was observed upon this occurring.

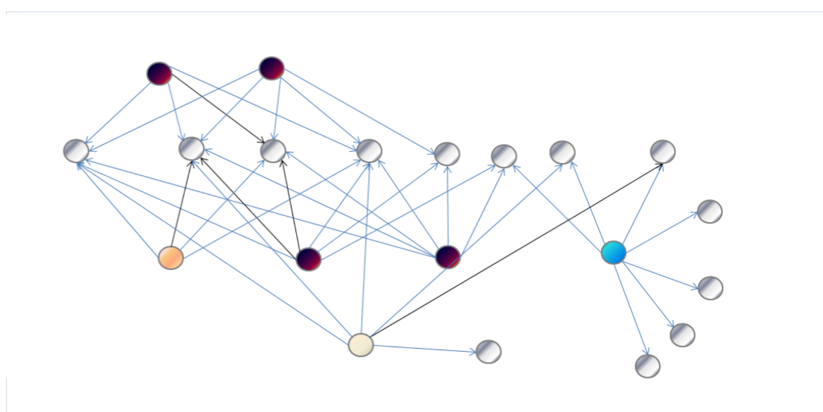

Figure 5.11: Illustration of trade network after bankruptcy of victim

Figure 5.11 shows the network that was formed after the bankruptcy of our central agent. The dark coloured nodes are agents who were both dependent on the victim, and were bankrupted some time after his bankruptcy. The blue node reflects an agent who was dependent on the victim, however, survived his bankruptcy. The peach node reflects an agent who was not dependent on the victim, but became insolvent after the victim. All other nodes survived and had no dependencies on the victim and they are either parent nodes (grey) or children (beige). Finally, edges drawn in blue are new, and edges drawn in black existed prior to the bankruptcy of the victim. For clarity, a key is provided in Table 5.2.

Shade	Meaning
Dark	Agent bankrupted, was dependent on victim
Blue	Agent survived, was dependent on victim
Peach	Agent bankrupted, was dependent on victim
Other	Agent survived, not dependent on victim

Table 5.2: Key for the colour coding of the networks

Firstly it is important to notice that the overall structure of the network changed dramatically in response to the removal of the victim. His dependents found new suppliers of Good 1, as illustrated by the many blue outbound arcs from the dependents in the diagram. Also notice, that the all the bankruptcies are on the agents who depend heavily on just 6 other agents. For clarity, this portion of the network is separated and shown in Figure **??**. Further, this figure shows agents who, prior to the bankruptcy were dependents themselves, or not dependent on anyone with noone dependent on them. They were the agents less good for intiating trade with and as such were self sufficient. However, on removal of the central node, these agents have a chance to become wealthy through new demand for their products. However, evidently they were unable to supply the quantity needed at reasonable price to all those agents who had been dependent on the victim. As such, the network has become overloaded. Even an agent who was not dependent directly on the bankrupted agent falls victim to the supply shortages and price inflation caused by this change in structure. It is clear that the model allows for agents to build new connections, interact with new agents when faced with this sort of economic crisis. However, in spite of this, they are still unable to survive and the reasons as to why are really quite simple. These new arcs, or dependencies that form were not apparent at meaningful levels of purchasing previously for a very good reason. They were not as good as the victims prices, they could not offer the victims quantities to all the agents. In essence, a new network structure forms with new dependencies, but these are not sustainable relationships. These nodes cannot cope with this new load and can't serve all the agents at the prices they need. As their stock pile of Good 1 is depleted, prices soar, and as a result become unaffordable to many. This explains why it is plausible that this extra load actually indirectly effects (not through a chain as such but through supply shortage and inflation) an agent who was in no way dependent upon the victim.

What is extremely interesting is those who do not succombe to bankruptcy and survive, even purchasing as high a quantity of Good 1, for the duration of the simulation. Again, for clarity, a separate illustration can be seen in Figure 5.13. Now, the brown node wasn't dependent on the victim in the first place, however he is heavily integrated in to the overloaded network. There is however a massive difference. He is in addtion dependent on a node outside of these Good 1 providing hubs. He is also a suitable trade partner for an agent with one two dependents. Two dependents is surely far more sustainable than six!

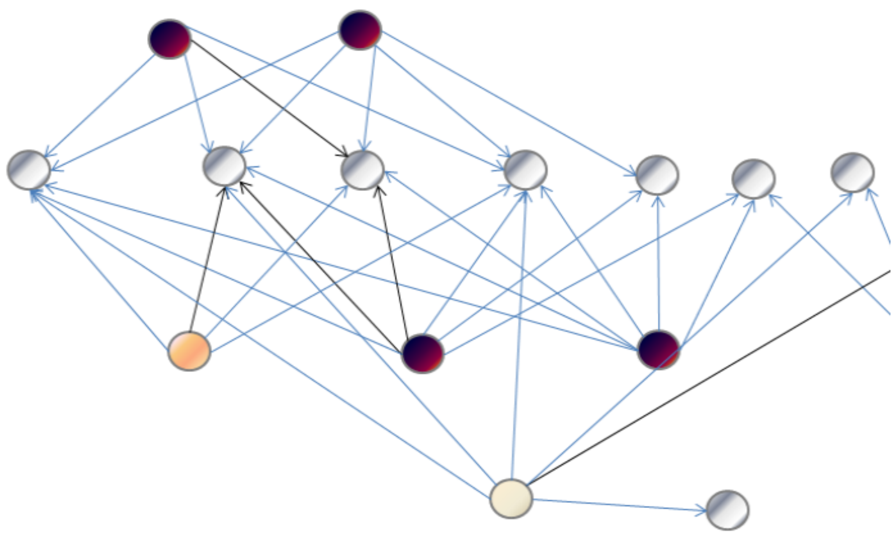

Figure 5.12: Illustration of common dependency

The blue node illustrates this point further. Previously dependent on our victim, he then forms an entirely different network virtually completely disjoint from the overloaded one. These agents provide only to him, and as such he is able to avoid bankruptcy through sourcing his Good 1 from an isolated set of agents, resulting in a reliable and sustainable network with redundancy in the case that one of the nodes he is dependent on suffers high prices or supply shortages.

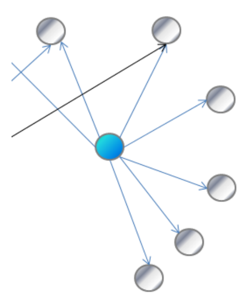

Figure 5.13: Illustration of survivors

We have seen that although the path lengths in the network prohibit the notion of bankruptcy chains effecting more that one person, the bankrupting of the central agent does have an impact on the ability of dependent agents to survive. We also saw that a change in network structure causes knock on effect to agents who weren't dependent on the bankrupted agent. It seems that dependency cannot capture, but can illustrate, the potential economic difficulties caused by removing an important node. In addition, it has revealed that hubs, or agents with high degree, are a source of vulnerability for its dependents when the dependent only engages in trade with other hubs. Low centrality means resources aren't stretched as far. It is evidently important to be dependent on not only central hubs, but also on nodes with low centrality. This avoids competition from other agents, and this competition contributes to supply shortages and price rises.

5.3.3 Conclusion

In order to anticipate the agents which will be bankrupted after, it is necessary to consider both the network prior to and after the the victim has been forced out of the economy. The agents likely to be bankrupted are those who were meaningfully dependent on the victim and who are attracted to the same set of agents, which become hubs as the network restructures. If these agents weren't previously hubs, or had considerably lower centrality, it is likely that the new load will not be coped with. In this situation all agents dependent on hubs and only hubs will be bankrupted. In the simulations studied, 75% of the time the hubs that emerged after the bankruptcy had previously had no inbound edges. However, after the bankruptcy, they had the highest centrality in the network. This shift illustrates the fact that previously these had not been the best trade partners and hence it is unsurprising that they could not cope with the new load. In addition, 75% of the time, the new hubs were actually agents who both had no dependents and depended on noone prior to the victim exiting the economy. This shows they were able to produce enough Good 1 to sustain them selves, but their stock piles didn't suit those in need of Good 1 as well as the victim.

This dependency network was useful in analysing the disruption caused by the most central agent exiting the network. The way in which there were hubs and these hubs were not linked to each other shows the network was composed in a similar manner to the star network - in which there is one central node which all nodes must go to. The network seen differed in that there were multiple hubs or stars, however the lack of transitivity, cycles and so on made for a fairly inefficient network, which is likely to contribute to the poor resilience of the dependents on the victim. To formalise as opposed to speculate the low efficiency on the network structure, and the vulnerability of the network to removal of a victim, the network efficiency can be computed as:

$$E(G) = \frac{\sum_{i=1}^{n} \sum_{i>j}^{N} \frac{1}{d_{ij}}}{N-1}$$

where E is the efficiency, G is the graph, N is the number of nodes, and $d_i j$ is the distance between node i and node j. If j is not reachable from i, the distance is undefined [14].

The average efficiency in the local disconnected network was 0.152 prior to the bankruptcy. The expected change in efficiency, or a measure of vulerability stemming from the deactivation of the victim v can be computed as:

$$C_v^I = \frac{\Delta E}{E} = \frac{E(G) - E(G')}{E(G)}$$

where *G'* is the network attained through removing the edges and node of the victim [14].

The average efficiency of the initial trade network was calculated to be 0.152 (values range between 0 and 1, 1 being the most efficient). After removing the victim node, the new efficiency was calculated to have fallen by 39%. When comparing this fall in efficiency to the actual computed efficiency of the new network that formed, it differed by only 3% on average. Thus despite the new formation of nodes, and the different nodes coming in to the equation, the formula was extremely accurate.

This fall in efficiency by over one third is indicative of the fragility of the network caused by such hubs and the short path lengths. Although it can recover structurally, unfortunately this isn't sustainable. In order to determine just how star like this is and thus how expected the bankruptcies are due to network structure, the trade network was compared to that of a star network of equal size. A star networks vulnerability comes from the single point of failure. The central node or hub is the only means of communication, or in this case, the only means of Good 1 reaching the other nodes or dependent agents. If the network formed is similar to that of a star structurally, it is reasonable to anticipate the catastrophic effect of forced bankruptcy on dependents, and thus can begin to analyse precisely why this structure amongst agents occurs. If however, it is not the case, then the contributing factor is not to do with a problem of network structure.

The similarity of the structure to that of a star network can be calculated as:

$$C = \frac{\sum_{i=1}^{g} \left(C_{max}^d - C_D(n_i) \right)}{(g-1)(g-2)}$$

where C is the centralisation index, C_{max}^d represent the actual maximum degree observed, and g is the number of nodes.

In fact, this indicator reveals that the network as a whole is fairly decentralised, with a centralisation index of on average 0.3 prior and 0.16 after the bankruptcy of the victim. This not only shows low centralisation to begin with, but in addition, it becomes more decentralised after the readjustment. This is in comparison to one in which a victim isn't chosen, where centralisation remains consistent throughout the simulation. The decrease in centralisation is perhaps quite telling of the struggles the agents face. Decentralisation occurs when more nodes are connected to other nodes - or there are more centers of gravity. However, this is indicative of the fact that a more centralised network is no longer possible after the bankruptcy. Agents used to be able to rely on a few select sources, however, removal of the important node forces agents to acquire goods from more sources, making the network more decentralised.

There is another problem in the results of this investigation which provides insight into the limitations of the model employed. The maximum path length in the trade network for a district is consistently one when conumption is enabled. There is no need for agents to buy from anywhere other than the producer. Thus no value is added in the traders purchasing a stock of goods as they have to initiate the next sale and no one buys it from them. Traders lose out when consumption is enabled since nobody depends on them. This is because if you need to purchase something from a trader, it is highly probable that you will receive both a better price and quantity from the producer directly. Since you have access to the producer, you will go there. In reality, abstractly, value is created when goods are moved from one location to another. The reason for the value being added can be anything from scarcity of the product where they are moved to, to the fact they are converted to a good which has more demand and thus higher value. For a more intuitive example, think about chocolate. Cocoa beans sourced from Africa are bought for an extremely low price from third world farmers. When they are sold to Cadbury's, a massive margin is made. Cadbury in turn makes its chocolate bars and sells these to Tesco for far more than the ingredients cost to purchase. Tesco sells these to consumers, again with a large mark up. The model however, doesn't reward traders well enough for this value creation, or in anyway at all. They buy cheap in one round and sell dear in the next, however, the fact that they are poor means not a lot gets moved and thus only a small margin is made unlike the producers who experience free exchange. In addition, their poverty increases what they charge and thus are never a valuable trade partner. This in turn means producers won't seek out traders as they can rely on traders seeking out them. This means that long trade chains aren't witnessed as everybody in need of purchasing a good can go straight to the producers. There are no middle men involved in the process. In reality, long trade chains exist, with each sale creating more value as a good is moved from one place to the next. In addtion, in the real world traders are not poor. On the contrary, merchants' knowledge of the trade networks they were involved in allowed exploitation of prices and they were some of the wealthiest. Even today, shipping companies make vast amounts of money - far more than the third world farmers. It is likely that the maximum trade length is a limitation. With consumption, the possibility of traders selling Good 1 is slim due to the poor prices they offer and the fact that they need it to survive. This topic is to be discussed in the final conclusion where suggestions for enhancing the model will be presented.

Chapter 6

Permitting Agents to Remember Encounters

So far, agents search through their district in order to determine who the best trade partner is. We saw in Chapter 3 the effect of restricting the number of agents that an agent is permitted to search through in order to find a trade partner. Now we ask, what happens if the number of people an agent can search through is restricted, but they can also choose to store those agents with whom they had an exchange, and search through agents in their memory in each round? For memory to be applicable or useful to agents, in the simulations used in evaluation trade is increased by raising the percentage of agents who can produce only one good to 20%.

An investigation into the usefulness of memory and also the effect it has on the number of distinct trade partners an agent has and the long term trading relationships that may emerge as a result is conducted. In addition, the effect of memory size is assessed to help determine the importance of memory to agents. Implementation of memory is analogous to sustained trading relationships. It is a way of agents being able to learn who is a reliable source of a specific good, and keep returning to that agent.

6.1 Motivation

The motivation behind remembering trade partners is twofold. Firstly, it is more realistic that an agent strategically records those who are suitable trade partners for them. In reality, we do not randomly select where we will buy our shopping, but we know from experience where the best places to go are. We know the corner shop is overpriced from having to pay 3 for tinfoil, and so only go there when we have no choice. We know that Tesco is quite cheap from having bought Tesco Value baked beans at 9p, but that, at the moment, if we want sausages, Sainsbury's has a discount. Although more factors than price come into it in the real world, and this is a massively simplified version of reality, the point is that we learn from experience where to shop. If an agent knows that he always trades with a few select partners, why not be sure of engaging in hypothetical trade with them in every round? Likewise, why waste time searching through contacts with whom he has never engaged in trade? Of course, agents will still use contacts, but a larger share of their potential partners will come from memory.

The second motivation for implementing memory was the introduction of learning, which is the topic of the next chapter. For this extension agents had to be able to remember agents that they learn about and hence it seemed important to also look in to the effect of memory in isolation from learning.

6.2 Implementation

In this section the logic involved in an agent making a decision as to whether or not to store (if it isn't stored already) an agent with whom he has engaged in trade will be covered. It is also necessary, since agents do not have a memory of infinite size, to explain the logic of deciding whom to replace in the memory in the case of an agent wanting to store a contact and his memory being full.

There are, therefore, two decisions:

1. Should this contact be stored?
2. If the agents memory is full, who should be replaced?

It may seem initially that the first decision is easy: if there is space, store the contact. However, it is not that simple. Since the majority of an agents potential trade partners will now come from its memory, and anticipating the implementation of learning, this means that you can't store just any contact. It is feasible to only have one exceptional contact who facilitates massive increases in utility among the ten contacts in an agents list of potential partners, meaning that even when an agent has free space in his memory, it is not desirable to add everybody with whom he is engaged in trade.

Let us consider what determines why one agent would want to store another. Firstly, it is relative in the sense that an agent is comparing the potential remembered agent to those who are already in his memory. Having considered what constitutes a worthwhile interaction, it seems that the following two points influence the decision.

1. Change in utility that the trade gave you
2. The MRS of the agent to be stored

While one might consider that change in utility and MRS will show the same thing so perhaps one could be used over the other? The problem is that the MRS of an agent captures the ratio of goods it owns in other words, how skewed its stockpiles are. By comparing the MRSes of two agents, their suitability - how well their stockpiles complement each other in trade becomes apparent.

Should one agent have an abundance of Good 1, and only a small amount of Good 2, while another has an abundance of Good 2 and only a small amount of Good 1, They would be suited since the one needs Good 1 to balance out his stocks, and the other needs Good 2. They could cooperate in trade so they both benefit from balancing out their stockpiles of goods. However, MRS does not capture the magnitude of the trade - the quantity exchanged. Utility, however, does. The change in utility shows how close the ones stock of Good 1 is to the others stock of Good 2 relative to the price. This is because if their MRS's were different, but the one was far less wealthy in how much of Good 2 he possessed, the other may not be able to sell him as much Good 1 as he would like. Thus, the change in utility indicates the quantity exchanged, which is an important factor in suitability.

Now that the contributors are known, it is possible to decide whether an agent is worth storing. Since, as mentioned, an agents value must be considered relative to the value of the agents already stored, we can assign each agent in memory a value reflecting its usefulness based on the two attributes, change in utility and MRS relative to the other agents in the memory. However, we have a twist in that *both* attributes must be taken into consideration. To do this, I adopted a notion from Decision Analysis that allows a value to be assigned reflecting the aggregate benefit of something relative to others. This aggregate benefit is a value, between 0 and 1 in this case, that encapsulates the value of every attribute to be considered. The steps taken are outlined below.

1. Calculate the aggregate benefit value of the agent to be considered to store, call this v_p
2. Calculate the aggregate benefit value of each agent in your memory, using the best trade as the trade for change in utility, call this set R_v
3. Compare v_p to each element r_v in the set R_v:
   (a) If there is an $r_v < v_p$ and there is space in memory, store the new contact
   (b) If there is no $r_v < v_p$, do not store the contact
   (c) If there is an $r_v < v_p$, and there is no available space, decide who to swap

Now we just need to see how an aggregate benefit value is arrived at. Let a_d be the agent making the decision, a_p is the potential agent being stored, and a_r^i is any remembered agent in the set of memory items, R_a. The idea is that a bigger change in utility gives a bigger value for utility, and a bigger difference in MRSes leads to a higher value for MRS. The values are then equally weighted (since their importances are equal) - they both carry a weight of one half. The values are multiplied by their weights and summed to find the aggregate benefit value. Formally they are computed as follows:

1. Find the range of changes in utility for all $a_r^i \in R_a$, and a_p. Let the minimum change in utility be u_m, and the range be u_r.

2. Find the range of absolute differences in MRS:

$$abs(a_d - a_i) \quad \forall i \in R_a, a_p$$

. Let the minimum difference be m_m and the range be m_r

3. For all a_r^i in R_a, and a_p, compute the value corresponding to their change in utility to be:

$$u_v^i = \frac{u^i - u_m}{u_r}$$

where
u^i is the change in utility generated for agent i, and u_v^i is the corresponding value for that change in utility

4. For all a_r^i in R_a, and a_p, compute the value corresponding to their difference in MRS from a_d to be:

$$m_v^i = \frac{m^i - m_m}{m_r}$$

where
m^i is the difference in MRS between a_i and a_d, and m_v^i is the corresponding value for that difference in MRS

5. For all a_r^i in R_a, and a_p, the aggregate benefit value for agent i, A_i is computed to be:

$$A_i = \frac{u_v^i + m_v^i}{2}$$

Now that the question of when to store has been answered, let us now tackle what happens if an agent decides to replace an agent in its memory in favour of a new one. How should who goes be computed? The method above could be used, but comparing agents in memory that you have already had exchanges with is more complicated. Issues arising include questions such as, how often are they used? How long ago was it that you had your best exchange? How does your most recent exchange compare with your best? In light of these questions, I decided to use the same technique of aggregate benefit value only this time with more attributes. You assign each agent in your memory a value reflecting its "goodness" relative to other agents in the memory. You then choose to substitute the new agent with the agent in memory with the lowest aggregate benefit value. This time, however, it is not the case that everything has equal weight, so instead we employ the method used in the *SMARTER* algorithm - Simple Multi-Attribute Rating Technique Exploiting Ranks. This algorithm assigns values to attributes and, by ranking the attributes, also assigns them weights. The weights are normalized, and the aggregate benefit value is the sum of the weighted value of each of the attributes. In fact, the above method computes values for attributes in the same way *SMARTER* does. Here, the alteration is in the way the attributes are weighted. All that needs to be decided is a ranking from most important to least important attribute. The attributes to be considered, in the order of importance are as follows[1]:

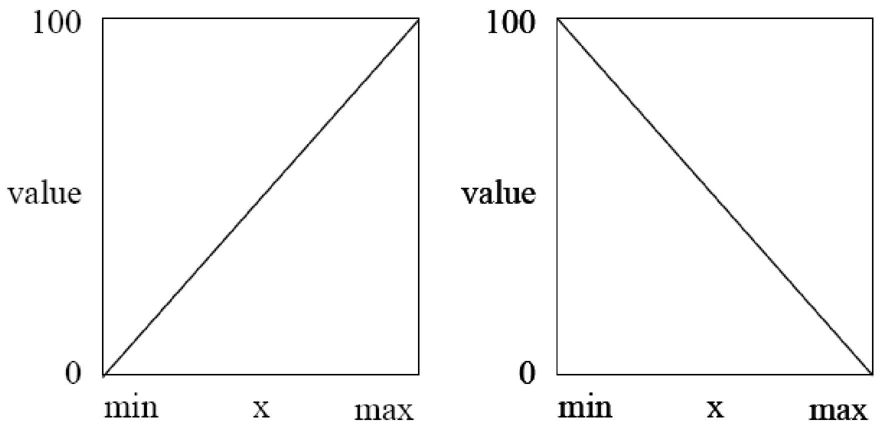

Figure 6.1: Graphs illustrating value functions for attributes

1. Change in utility
2. Difference in MRS
3. Number of uses relative to the time the contact was added
4. Difference between the change in utility of the most recent trade, and the change in utility of the best trade
5. Amount of time since the best trade occurred

The attribute values were computed using the same method as above, using an assumption that the relationship between an increase (or decrease) in the actual value of the attribute (for example the change in utility) was linearly proportional to the change in computed value of the attribute (the value of the change in utility relative to the other agents change in utility). Plainly, this means, for example, that there is no difference in the change in value if the change in utility is increased slightly from when it was initially low to when it was initially high. Graphically, this is shown in Figure 6.1[9].

Some attributes however, are better when they are low. For instance, it is better for there to only have been a short time that has passed since the best trade - for this the value is computed in the usual case, and is subtracted from one to represent this inverse relationship.

The weights were computed using then Rank Ordered Centroid (ROC) technique. For n attributes, ranked from 1 to n, the ROC weights are given by:

[1]The order was experimented with and judged based on the percentage of trades occurring with agents from their memory - the following order was found to be optimal

$$W_i = \frac{1}{n}\left(\sum_{j=i}^{n}\frac{1}{j}\right)$$

It is now necessary to investigate the results that this extension had on the evolution of the simulation, including any changes in agent behaviour and the level of trading.

6.3 Evaluation

In order to evaluate the effect that memory had on the evolution of the simulation, two things were examined:

1. Loyalty
2. Specialisation

In addition, the number of agents that an agent could search through to find a trade partner, which shall be termed "sight" was varied from 5 to 20 (20 being all agents in the district) with increments of 5. In order to fully see the affect of memory on the evolution of the simulation, only 20% of agents were able to produce both goods. Ten simulations were run for the topologies Local Connected, Local Disconnected and Small World Network for each of the values for sight, and for memory being enabled and disabled.

The reason for varying sight was to see if the "usefulness" of memory differed with the number of agents an agent could search through.

We have already seen the method for characterising the level of trade - the percentage of trade and of specialisation in the form of the continuum ranging from pure producers to pure traders. However, loyalty has yet to be mentioned. Loyalty is the idea of agents returning to the same trade partners. In order to measure loyalty, an adapted Herfindahl index was utilised- an idea from Wilhite's paper [22]. This index is actually used in measuring industrial concentration, but was adapted in order to reflect the information that we are seeking. The calculation is based on the number of times an agent engages in trade with the same partner. The loyalty for agent i can be calculated as:

$$L^i = \sum_{j=1}^{k}(100a_j)^2$$

where a_j is the proportion of all trades, initiated by agent i with agent j, and agent j is one of the k distinct agents with whom agent i trades. The maximum value L^i can take is 10 000, and this indicated maximum concentration. In other words agent i always trades with one agent. The index was calculated for each agent in the categories heavy producer, heavy trader, and pure trader. An average is taken for each of these categories, and this is averaged over the repeated simulations.

The level of trade was actually increased and decreased at different values of sight across the topologies. However, after performing a t-test for significance, the effect that memory had on the level of trade turned out to be insignificant for all sight values used. Thus no further discussion of the level of trade will be carried out.

However, there were significant differences in both specialisation and loyalty, and these will be the topic of the following two sections.

6.3.1 Loyalty

As previously mentioned, the average loyalty was calculated for each of the specialisation categories other than the pure producers. The reason for the exclusion of pure producers is straightforward - they rarely if ever trade and as such the index is misleading and unnecessary. The reason for comparing within specialisations and not across specialisations is that the number of trade partners differs significantly across specialisations and thus the index is inflated for those who trade less. However, the change within a specialisation is subject to less variation, therefore they shall be compared in an isolated fashion.

Heavy Producers

For heavy producers, the loyalty index is consistently and significantly greater with memory than without across all topologies. This indicates that the ability of an agent to store trade partners allows them to learn where the best deals for them are and as such they can continue to return to their favoured partners. In addition, as shown in the graph depicting the Small World Network in Figure 6.2, memory makes the most significant difference when sight is low. When sight is low and memory is disabled, an agent can only search through a subset of the population, chosen at random, in order to find a trade partner. Thus, if a good trade partner is found, the probability of picking them as a potential partner on the consecutive round is equal to the probability of picking any other contact, so potential trading opportunities can easily be missed. However, when memory is introduced, over the course of a few rounds an agent is able to "sample" a fairly large proportion, if not entirety, of the population. Hence, if on the first round he finds a trade partner and decides to store them, then the agent has the ability to return to this reliable source in the consecutive round. Therefore, loyalty is greater when sight is low since an agent is no longer randomly sampling the population for trade partners. Instead they are learning who the best partners are. This trend is apparent across each topology investigated.

An interesting point in the graph is the relatively sharp decline in loyalty between a sight of 5 and 10 when memory has been employed. This is perhaps due to the fact that this range in sight has the most significant impact on loyalty due to the difference made between being able to search through a quarter of and half of the population. Here, the increase in the number of potential trade partners on each round allows for more trade partners to be found. In turn,

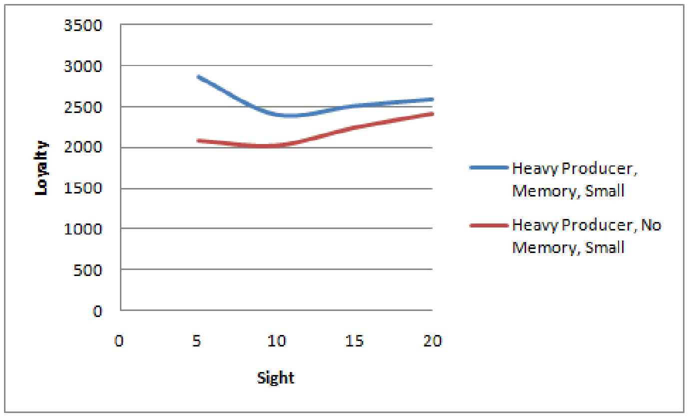

Figure 6.2: Loyalty in the Small World Network, with and without memory

this increase in sight allowing agents to find more trade partners outweighs the effect memory has on loyalty and this results in a decrease in the index.

Heavy Traders

The graphs for heavy traders actually have an intersection on each topology, where the loyalty when using memory falls under the loyalty when they are not. In each of the topologies this happens after a sight of 10. The reason for loyalty to be greater with memory is the same reasoning as for heavy producers. However, the intersection is significant in every topology except for the Small World Network. Perhaps it would be useful to establish the reason for the intersection at all prior to establishing the significance of the point at which they intersect. An illustration can be seen in 6.3

As we know, agents who trade more have higher numbers of partners. In addition, it is likely that these partners change as the simulation evolves. By restricting himself to searching through agents in his memory, perhaps an agent will miss out on opportunities to find new, long-lasting relationships. If the contacts in his memory are less useful or not offering favourable deals, and he is unable to sample the entire population of contacts (even though this is possible when sight is 20) then he may end up making deals that are not as good with the random sample from the population.

For example, imagine that an agent A knows a great source for Good 1, B, with whom he trades frequently and has stored in his memory. Then perhaps, since everybody in the population is able to search through more agents when sight increases, someone else, C, also finds that B is beneficial to them too - after all, the probability of someone else finding B is higher when sight is higher. Furthermore, C is actually also better suited to B than A. Now, the majority of the time, when A engages in hypothetical trade with B, he finds that C must have got there first. Bs goods are depleted and even, and A and

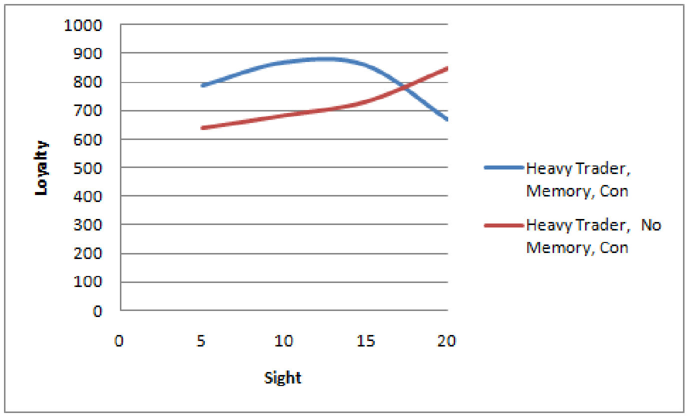

Figure 6.3: Loyalty in the Local Connected Network, with and without memory, for heavy traders

B are no longer good trade partners. With no memory, this is not a problem. A can go on to find a new long term trade partner. With memory, however, B is still in As memory and it is likely that most rounds A will attempt to trade with B. Of course, there will be the occasional round that A gets to B before C, but not as often. Thus, A must either produce, or find a random and potentially unreliable agent to trade with. This in turn means that his loyalty may decrease as sight moves from 15 to 20 due to more competition for trade partners, which causes agents in his memory to "expire" when they are used less frequently. Since the majority of As contacts come from memory, and the randomly selected agents aren't rechosen if they are in your memory (since this would defeat the point of the stochastic element when sight is high; it would not truly reflect how useful memory is, as it would behave in the same manner as sight) he is potentially left with an out-of-date memory and is hindered by the randomness in selecting trade partners.

Now the reasoning is established, we are in a better position to consider why the intersects differ in the Local Connected and Local Disconnected Network. In light of the explanation, the differing intersects are fairly trivial. In a Local Disconnected Network, competition is far higher, and closed borders means no goods flow in or out to open up new trading opportunities. Thus sight does not have to be as high for this problem to present itself. The graph illustrating the intersect in the Local Disconnected Network can be seen in Figure 6.4. The opposite applies for both the Local Connected and Small World Networks, which displayed very similar trends.

6.3.2 Pure Traders

The Local Disconnected Network showed by far the most significant difference between loyalty with and without memory for traders (due in part to the negli-

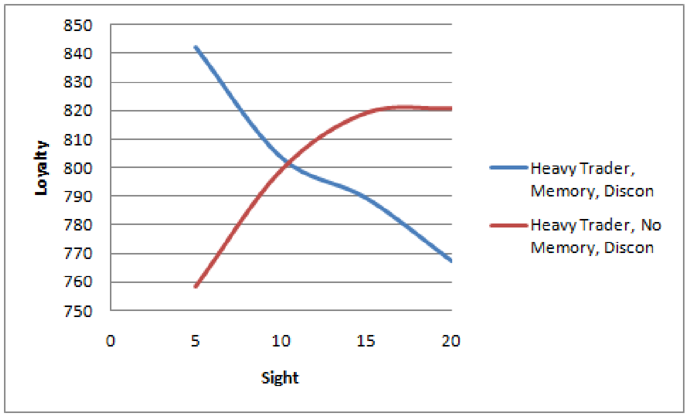

Figure 6.4: Loyalty in the local disconnected network, with and without memory, for heavy traders

gible variance across simulations). However, all topologies showed a significant difference in loyalty. For the Local Disconnected and Small World Network, loyalty of the pure traders with memory was constantly above loyalty without memory. The graph of loyalty for pure traders in a Local Disconnected and Small World Network can be seen in Figures 6.5 and 6.6 respectively.

In the Local Disconnected Network, the convergence between a sight of 15 and 20 again occurred prominently between the two lines. In contrast, this similarity was apparent in the Small World Network when sight was between 5 and 10, it widened between 10 and 15, before converging again between 15 and 20. Although convergence in the 15 - 20 region can be put down to the same reason laid out for heavy traders, the closeness of the lines between 5 and 10 cannot be attributed to this.

Perhaps the importance of trade for a pure trader is so high that, irrespective of whether or not memory is or is not employed, when sight is low, the benefit of increasing it will lead to an increase in loyalty of an equal rate. This is due to the fact that the extra agents that can be searched through allows the discovery of good, reliable trade partners, and thus sight outweighs the added benefit of memory. However, this is apparent in the Small World Network and not the Disconnected Network. This is most likely due to the fact that, again, the loyalty in the Disconnected Network suffers since closed borders mean that there are no inflows and outflows of goods, which in turn means that trading opportunities through direct or successive trades prevent the equal number of changes in stockpiles of goods. This results in each agent having fewer potential trade partners (they are exhausted more quickly due to the lack of diversity.). However, pure traders almost always find trade more beneficial than production, so random unreliable partners are better than none. When sight is low, and the subset of the population to which an agent is suited to trade with in the long term is small, the probability of choosing one of these agents without memory is

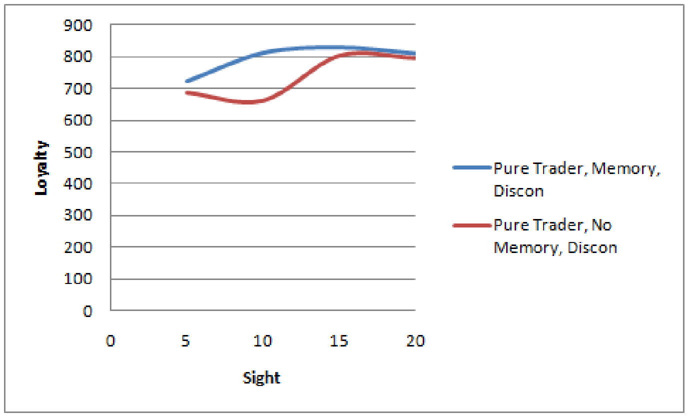

Figure 6.5: Loyalty in the local disconnected network, with and without memory, for heavy traders

in turn low. The Small World Network seems far less affected by this - although loyalty with memory is higher, loyalty without memory increases at the same rate - it is simply capped by the fact that agents cannot learn.

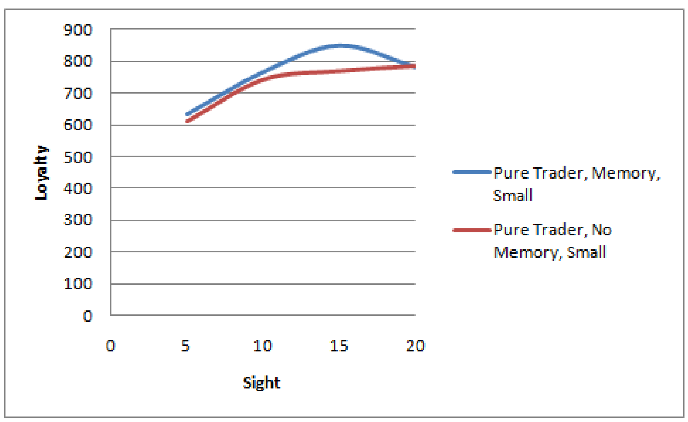

Figure 6.6: Loyalty in the Small World Network, with and without memory, for heavy traders

This trend of loyaly is indicative of the difference that memory makes to agents. It helps them to organise themselves to deal only with a small subset of the population and form long-lasting trade relationships with their peers. This self-organisation is an important and interesting emergent phenomenon. It better reflects reality - agents become more selective. In addition, important thresholds were found which again emphasise the importance of initial conditions on the evolution of the simulation.

Let us now move on to specialisation to see if varying sight and the addition

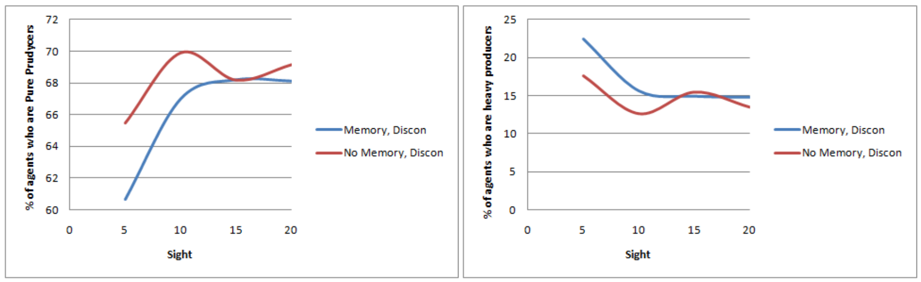

(a) Pure Producers (b) Heavy Producers

Figure 6.7: Specialisation with and without memory in a Disconnected Network

of memory alters agents decision making with respect to the interactions they most frequently perform.

6.3.3 Specialisation

Memory had a significant effect on the pure producer and heavy producer specialisations, but no significant effect on traders, most likely due to the fact that trading is a last resort, thus memory will not force any traders into production, nor a significant number of producers into trading.

In the pure producer category, the percentage of agents falling into this category was consistently lower with memory than without, again especially when sight was low for both the Local Connected and Disconnected topologies, Figures 6.8a and **??** respectively. Memory allowed for occasional trades - particularly at the beginning of the simulation when the rush of trades occurs - to be stored and reused. Thus more trades occurred, most likely moving them into the heavy producer category. In addition, as the sight was increased, the percentage of pure producers increased. On the face of it, this seems slightly counter-intuitive. However, more sight means more heavy producers can rely on being selected by other agents as trade partners thus do not need to trade themselves. This increase in pure producers occurred in each topology, both with and without memory, before plateauing or in some cases falling slightly as sight passed 15.

In the Small World Network, however, there was again an intersection between the two lines, where the number of pure producers with memory enabled exceeded the number with it disabled. One good explanation for this is that in the Small World Network, even more producers can rely on agents initiating trade with them when sight is high, due to the freer flow of goods revealing new trading opportunities. This bonus of the Small World Network is exaggerated with memory, as agents are not relying on random selection, but strategic selection, making it even easier for pure producers to depend on being found. Similarly, once again the Small World Network has an intersect at which heavy producers with memory fall below heavy producers without. This shows that

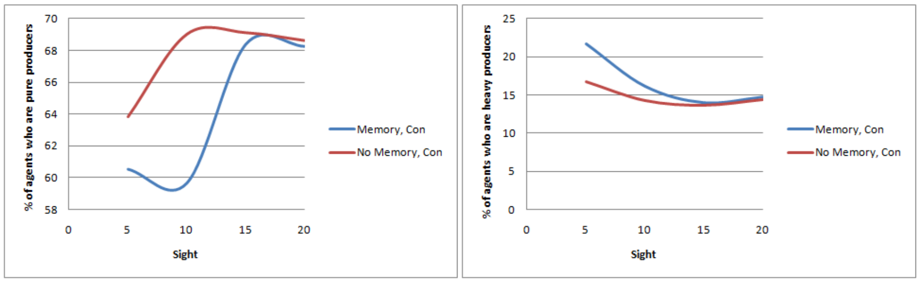

(a) Pure Producers (b) Heavy Producers

Figure 6.8: Specialisation with and without memory in the Connected Network

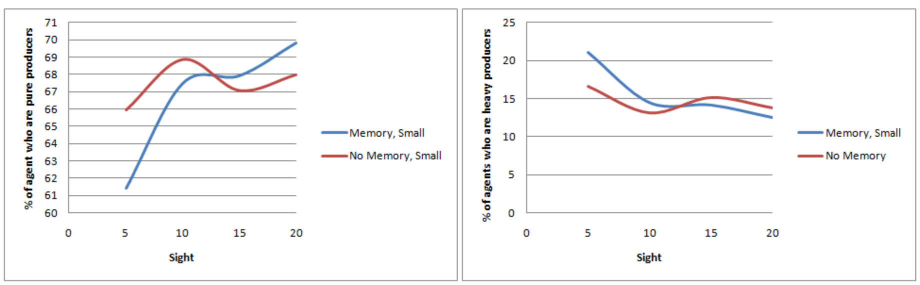

(a) Pure Producers (b) Heavy Producers

Figure 6.9: Specialisation with and without memory in the small world network network

the pure producers and heavy producers swap at this threshold value for sight of 12.

With respect to the heavy producer category in the Connected Network, the percentage of heavy producers was always higher with memory, although this difference became less significant again as sight was increased. This is due to the fact that the opportunity for strategic selection allowing for repeated trade made meant that some pure producers were likely to now fall into the heavy producer category as the occasional trade they made had a higher probability of being repeated a sufficient number times for them to pass into the heavy producer category. However, just as the number of pure producers increases with sight, the number of heavy producers falls. The can be attributed to the tension between repeating trades being possible leading to pure producers becoming heavy producers, and agents who can search through more agents meaning heavy producers can rely on "free exchange" and becoming pure producers.

In the Disconnected Network we also witness the repetitive occurrence of sight being 15as an interesting value. For heavy producers, it is the only time in which the percentage of heavy producers when memory is not employed is greater than that when it is. This further indicates the importance of this value of sight in the Disconnected and Connected network. At this point, there are some agents who are pure producers and some who are heavy traders. There is obviously considerable tension between the use of memory and the magnitude of sight that takes effect at this point in both the Local Connected and Disconneced Network with respect to specialisation. The tipping value for the Small World Network, however, is lower. Here sight does not have to be as high in order for specialisation to be affected, i.e.for agents to migrate to heavy trade. This can be attributed to the trading opportunities that exist in the Small World Network but not in the other two. Thus, agents have to be able to search through less of the population before heavy producers can begin to rely on the demand of other agents.

6.4 Conclusion

The varying of sight caused tensions between being able to rely on "free exchange" and more trade, and thus tensions between specialisation breakdowns. However, there were some clear threshold values witnessed.

Self organisation was witnessed as agents began to trade with increasingly smaller sets of the population, but the downside of memory presented itself in the form of a constantly changing world leading to some agents being stuck with artefacts of exchanges that had now been taken from them as increases in sight lead to increases in competition.

The increase in competition, however, was realistic. The more global the districts became, the more likely it was that someone could be a better trade partner to a producer. This is akin to the real world. The more choice people have, the more competitive the market is and the majority of the time in business, as in the simulation, the best price wins.

It is thus not surprising that the pure traders of the Small World Network using memory suffered the biggest hit as sight was increased from 15 to 20. Opportunities were stolen and their loyalty declined, indicative of less reliable trade partners due to competition and artefacts in memory as the producers found better buyers.

Chapter 7

Learning: Trade of Knowledge

So far the effect of the initial topology on the evolution of the simulation has been examined, and the agents have been seen to specialise in with whom they engage in trade. However the evolution of trade networks has not yet been permitted. For this, a new sort of interaction was implemented - learning. In the simulation learning constitutes being made aware of another agent, preferably a "foreigner" with whom it may be beneficial to engage in trade. An existing contact in your district or memory is chosen from whom an agent can learn. In a *quid pro quo* world, the agent gives up a contact beneficial to the learner, and the learner gives one to the agent. It is preferably, but not necessarily a foreign agent that is traded, since in the beginning it would be unlikely for an agent to have knowledge of a foreigner that their learning partner doesn't know of. As the simulation progresses however, the knowledge of the existence of agents propagates through the network leading to new opportunities and new sufferings. Profits are made as reliable agents gain a global reputation, and *creative destruction* is witnessed in this increasingly competitive economy.

7.1 Motivation

As mentioned, learning provides a way for trade networks to evolve. In essence, it permits globalisation to evolve, employing the most efficient agents to emerge as new trading links to the rest of the world. But we already have trading links to the rest of world! Why is it important for these to be dynamic and adaptive? At the moment, trading links across districts can only exist if there is a crossover agent linking one district to another. Although this allows goods to flow around the network, and through a series of trades can reach any agent in the world, the crossover agents are static. Thus, they are blessed with this advantage at the beginning of the simulation, and this remains the case throughout the entirety of the simulation. However, the crossover agent could be poorly endowed with respect to its ability to produce. It could in fact a poor agent who provides no use as a trade partner to anyone in his home or foreign district. Perhaps he is a pure producer, the kind who noone cares to initiate trade with. In these cases, he is blessed, for no reason other than chance, with a strategic position which he, and noone else, is able to exploit. Surely it should be the case that those who act as links to other districts are the best link possible. They offer the best prices, the best quantities, they are the importers (purchase Good 1) and exporters (sell Good 1) of the economy.

Learning allows agents to learn of an agent that may be of use to them. If an agent is given a great source for Good 1, and his similar friend wants a trade partner to acquire Good 1, then learning means that he can get it. Now this great source of Good 1 is in demand by two extra agents. Those who should be links, become links. It is no longer a random endowment. This means that as goods spread globally around the network, agents evolve to be importers and exporters for their district, and true efficiency in globalisation is the result. Perhaps we could then identify more attributes of the network such as who the most important agents are or which districts are the most important.

Let us consider why this is of interest in Economics. In honesty, it does not get as much focus as one might think - but that is not to say it shouldn't. The debate over the effect of globalisation is ongoing. The aim is not to answer this question, but instead to assess whether or not the simulation can evolve in a way analogous to the real world, and to assess the validity of studying network topology and evolution as a way of investigating the effect of globalisation.

7.1.1 The Importance of Structure

In the background section it was briefly mentioned that economics is focussed very much on elements of a system as opposed to the connections between them. This is also true when it comes to the topic of globalisation. Economic globalisation can be seen as the act of integrating local national markets leading to the emergence of a global market place.Globalisation shouldn't be thought to be a purely economic context. It is equally related to social, cultural, pollitical and even technological emergence. However, herein, globalisation refers to solely economic globalisation as the simulation does not capture any other concept. Economics provides measures for assessing globalisation, or how *integrated* a country is into the global market. Generally these measures include imports and exports as a proportion of national income, the weight of how much you export relative to imports, immigration rates and the extent to which foreign technology is used. Ironically, these are also in economics known as *flows*. These may seem perfectly viable measures of globalisation. However, whether or not this can really tell us much about the nature globalisation is debatable. These measures are calculated on a per country basis. It seems where these flows go to and come from is largely ignored. Even if this break down is provided by a country, it tells us little in isolation.

Surely it is necessary to examine the "big picture". Where do these exports go, where do they come from, and how are all the countries interconnected?

Surely these questions would better enable us to assess the effects of globalisation since seeing globalisation from a global view will allow us to analyze dependencies and true macro trends. This is likely to shed light on questions such as why has globalisation not benefited everyone? How has the “global view” evolved over the past 20 years? Have political issues caused structural changes in globalisation? Are emerging economies truly becoming integrated globally - are they becoming central in the global network? The study of networks seems to be invaluable here. If we can begin to understand the overall structure and evolution of the so called global network, we can not only be better equipped to answer these questions, but also assess whether countries , if any, are in fact truly global. We can even begin to figure out how the network responds to changes in political climates, how countries can become better integrated, and perhaps even why some poor countries are still so poor.

Having established the applicability of networks to the assessment of globalisation, it is necessary to assess what can be used to facilitate modelling such a complex network. In a paper written on the architecture of globalisation, [12] some answers that could bode well with the simulation employed were presented. This paper modelled flows of imports and exports. Countries acted as nodes and arcs between them represented trade. Both exports and imports were modelled this way for the years 1992 and 1998. The model was further enriched by altering the threshold for there to be an arc between two nodes. Thus, instead of having an edge between two nodes if there is any level of trade, an edge is only apparent if the amount of trade is greater than 2% of the countries total exports for instance. By creating this network, it was possible to assess nodes who were central to the network, or hubs, and trends in the overall network topology. This can be applied very nicely in the context of the simulation - there are districts which can be act as countries, and imports and exports constitute cross-district trade. Therefore the necessary data exists, enabling an investigation into whether or not the evolution in the simulations is in any way similar to that of our world to be conducted. In addition, whether there are correlations between wealth and important nodes in the network will be investigated, along with whether the results are in line with the Neoclassical view of globalisation. By doing so, it will become apparent just how applicable agent based modelling is to the field of economics if a new perspective from complex systems is taken on, and in addition what it can tell us about the world we live in.

7.2 Implementation

An agent would choose an agent to learn from and they would exchange an agent that they both (if possible) did not already know. In that round, if an agent chooses to learn, they are not permitted to engage in trade or production since learning should not be free, it should take time. An important thing to note is that learning is another form of *exchange*. The learner does not just receive a useful contact for free, they too have to share their knowledge of useful agents.

This notion of reciprocal learning came from a paper on knowledge diffusion through networks. [13] This paper suggests that a method whereby knowledge is given away is not supportable when there is a possibility of any secondary competitive effects that may arise from the free distribution of knowledge. In the context of the simulation, these competitive effects are possible. For instance, let an agent, call him a_l be a learner, and the agent from who he is learning be a_i. In a gift giving world, let the agent given by a_i to a_l be a_g. It is likely that a_l and a_i would have similar trade partners in terms of suitability - this is why a_l chose to learn from a_i after all. Therefore it is possible that the best agent for a_l corresponds to an agent that a_i already trades with. By a_i simply giving a_l this new trade partner a_g, it is thus possible that now a_g now prefers to trade with a_l and hence a_i loses! However, if an agent is exchanged, although the loss may still happen, it is more acceptable to a_i as he gained something in return. Hence the diffusion of knowledge of agents should be on a *quid pro quo basis*.

Learning was implemented through a Python script responsible for the following:

- Should the agent learn?
- If so who should they learn from?
- What agents should they exchange?
- Should the two agents store the agent they were given?

7.2.1 Decision to Learn

For the first question of whether or not an agent should learn, it was necessary to determine what contributors we have. In addition, instead of making learning a deterministic action, since it is difficult to quantify the benefit, it was implemented such that depending on your situation the *probability* of engaging in learning would vary. Learning facilitates an agent acquiring new agents with whom he is suited to trade. Therefore if an agent is a pure producer, it is unlikely that expanding his knowledge of who he can trade with is of much use to him. Therefore the probability of learning should be fairly low. If however, we are dealing with a pure trader who relies on good deals and reliable agents for survival in the economy, learning should be something he is willing to give up time to engage in.

Secondly we have the amount of time that has passed since his best trade. If he has not achieved a good trade for a long time, perhaps this is indicative of a changing environment, and perhaps he must branch out to find new opportunities. This therefore also increases the probability of learning.

Finally we have, in the case of consumption being used, whether or not an agent can afford to learn. If learning would cause the agent to be unable to consume the baseline quantity of goods, which would cause him to exit the economy, he definitely should not engage in learning since it would constitute a sort of economic suicide. For this attribute a binary variable was employed. If an agent can afford the time given up in learning, it is set to 1, otherwise it is set to 0, and consequently the probability of learning becomes 0.

A value is given to each of these contributors and they are weighted and combined as in memory - using aggregate benefit values, and the resulting value is the probability that an agent should engage in learning. A random number is generated uniformly between 0 and 1. If the number is less than or equal to the probability of learning, the agent learns, otherwise it doesn't, and instead goes on to choose between production and trade.

The final calculation is:[1]:

$$P(learning) = a \times ((0.2 \times b_v) + (0.4 \times t_v))$$

where

a is the binary variable of being able to afford to learn, t_v is the proportion of time spent trading, *0.2, 0.4* are the weights assigned to the corresponding attributes and b_v is calculated as

$$1 - b_t/t$$

where
b_t is the time of the best trade, and t is the current time.

7.2.2 Deciding who to learn from and who should be exchanged

Upon making the decision that it could be beneficial to learn, it is necessary to decide who you will source information from. An agent can source information from any of their contacts, either permanent or stored in memory. The agent searches through the foreign contacts of their contacts. It calculates the value of each contact that it could learn of. The one with the highest value is taken and in return, the agent finds a good contact to give to the agent it learned from. Again the technique of aggregate benefit values is employed since there are multiple attributes that influence the decision:

- How different your Mr's are relative to the other possibilities
- How good the price of exchange is relative to the others
- How close your stock of goods are in terms of compatibility and trading an optimal quantity

Let a_l be the agent learning, and a_c the agent who you may choose to learn of. The values for the attributes above are calculated as follows. For *MRS*, you are most suited to trade with agents with a very different *MRS*, so the value

[1] The constants were found through experimentation to make sure learning didn't damage the wealth of agents severely but still had an effect on the simulation

of difference in *MRS* is high if the difference is big, and low if the difference is small. Formally,

$$m_v^c = \frac{m^c - m_m}{m_r}$$

where
m_v^c is the value corresponding to that difference in *MRS* relative to the other agents, m^c is the difference between the *MRS* of a_l and a_c, m_m is the minimum difference in *MRS* between all candidates, and m_r is the range of differences in *MRS* between all candidates.

With respect to the prices, if an agent is selling Good 1, he wants the price to be high, and if he is buying he wants the price to be low. Since price is calculated as:

$$P = \frac{g_2^i + g_2^j}{g_1^i + g_1^j}$$

and *MRS* for agent i is calculated as:

$$MRS^i = \frac{g_2^i}{g_1^i}$$

where
g_1^i, g_1^j, g_2^i, g_2^j are the stock of Good 1 and Good 2 for agent i and j respectively.

then the price of a good always falls between the *MRS* of the 2 agents engaging in trade.

Thus giving a value to the price you would exchange at, relative to that of the other candidates can be computed as follows:

1. If a_l would be the buyer,

$$P_v = 1 - \frac{p(a_l, a_c) - m_m}{m_r}$$

2. Otherwise,

$$P_v = \frac{p(a_l, a_c) - m_m}{m_r}$$

where
P_v is the value corresponding to the price, $p(a_l,\ a_c)$ is the price of goods between a_l and a_c, m_m is the minimum *MRS* between candidates, and m_r is the range of *MRS* between candidates.

This means that the if an agent is selling, the higher the price relative to the range of possible prices, the higher the value. Similarly, if an agent is buying the lower the price relative to the possible range in prices, the higher the value.

Finally we have the value of the compatibility of your stock piles of goods. As mentioned before, it could be the case that the *MRS's* differ considerably between two agents, but one of them has considerably more to sell (more Good 1), or alternatively more purchasing power (more Good 2). This means that the quantity exchanged may be far less than optimal for one of the agents. They may have been hoping to buy or sell more. The only quantity you need to be concerned with is the quantity owned by the agent, of the good he would be giving away. Since it is not necessarily guaranteed that the price will approximate to 1, it is also necessary to consider the price. The value corresponding to your relative stock piles of goods is computed as follows:

1. If a_l would be buyer,

$$Q_v = \frac{q_1}{q_2}$$

   where

$$q_1 = \min\left(\frac{g_2^l}{p}, g_1^c\right) \quad q_2 = \max\left(\frac{g_2^l}{p}, g_1^c\right)$$

2. Otherwise,

$$Q_v = \frac{q_2}{q_1}$$

   where

$$q_1 = \min\left(\frac{g_2^c}{p}, g_1^l\right) \quad q_2 = \max\left(\frac{g_2^c}{p}, g_1^l\right)$$

and

$$g_1^c, g_2^c$$

are the stock piles of Good 1 and 2 respectively for the candidate agent a_c

$$g_1^l, g_2^l$$

are the stock piles of Good 1 and 2 respectively for the learner agent a_l
and p is the price between a_l and a_c

The final calculation for the aggregate value is the summation of the weighted values of the attributes,

$$V = \frac{3(M_v^c)}{5} + \frac{P_v + Q_v}{5}$$

The agent with the highest value is who the learner learns of, and the agent who knows this successful candidate is the learner's partner in exchange. The learner employs the same method to find who to give his partner. They both then decide independently whether or not to store the agent using the method explained in the memory section. However this time, the change in utility is not known, so instead it is calculated as a hypothetical trade - if I were to trade with this new learned agent, what would be my gain?

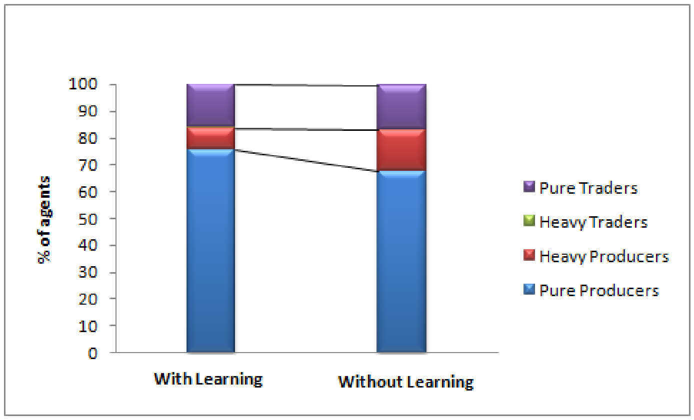

Figure 7.1: Chart illustrating effect of learning on strategy

7.3 Evaluation

The aim is to explore, on two levels, the effect of learning on network structure, and the evolution of trade networks. The first level is what it means for individual agents. What new patterns emerge? Are strategy shifts witnessed and do prices converge more readily? On the second level the focus is not on agents but on districts. Since districts are regions of isolated agents, they act as countries or clearly defined bordered regions. Trades between regions are analogous to imports and exports. Are there correlations between levels of imports and exports and wealth of districts? Do all districts interact with each other? Are there differences in the levels of imports and exports between districts? To do this, we will introduce a notion of networks based on imports and exports. This was applied to real countries by Raja Kali and Javier Reyes in a paper "The Architecture of Globalisation" and the idea of networks based on imports and exports, and comparison with my results to the real world comes from here.[12].

7.3.1 Specialisation, wealth and price dispersion

As stated previously, learning is an interaction to be performed in lieu of both production and trade, which accounts for the cost of networking. One would expect that this new interaction results in a decline in trade, since probabilisticly, agents who engage in trade more frequently engage in learning more frequently. In fact, the slight drop in the average level of trade, from an average of simulations using the local connected and small world networks actually resulted in an statistically insignificant fall in trading. Trade fell by 0.8%, and due to the variation in the percentage of trade across simulations, this was shown to be insignificant via a means test. Thus, learning cannot be said to either increase or decrease the level of trade in the simulation.

However, once again interesting strategy shifts are witnessed in the decisions

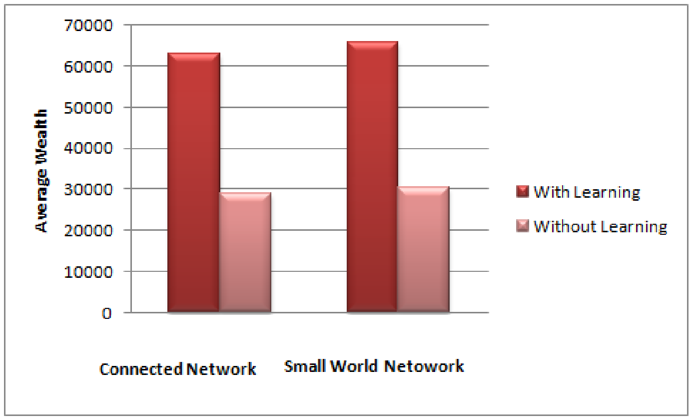

Figure 7.2: Chart illustrating effect of learning on average wealth

of agents. Figure 7.1 illustrates a somewhat counter intuitive alteration of strategy. It can be seen from the chart, that when learning is employed, there is a significant increase in the percentage of pure producers in the population. It is also apparent from inspecting this chart that there is a simultaneous drop in the percentage of heavy producers. The other two specialisations go virtually unaltered. From this, one can infer that with learning enabled, some heavy producers migrate to the pure producer category. This can be explained by the diffusion of knowledge of good trade partners through the network. As agents are able to learn on suitable trade partners, the make new previously unreachable contacts from foreign districts. Recall a heavy producer would produce for a few rounds and then trade in the next before returning to production. Now however, as an agent becomes aware that this heavy producer can offer a good trade, the heavy producer can rely upon trade being initiated by a pure trader. Since the trader is likely to store this good trade in his memory, he is also likely to return. Thus new long term relationships are formed and the heavy producer is rarely, if at all, required to initiate trade.

Another interesting and hugely significant effect that learning has on agents in the simulation is the effect of average wealth. Figure 7.2 illustrates a chart showing average wealth in both the local connected and small world network. It is clear from the chart that the addition of learning causes a large increase in average wealth of agents. In fact, it causes average wealth to double, with extremely low variation. This is quite remarkable, and the reason is not what one may expect. It could be thought that this new interaction provides a way for the traders of the world to make more profit and as such have a higher wealth, so the increase in wealth is likely to stem from the traders becoming wealthier. However, the distribution of wealth remains constant with and without learning. This implies that the increase in average wealth effects the population as a whole. Everybody gets richer, but nobody relatively richer. In other words, learning offers new efficiencies and opportunities in trade which lead to society

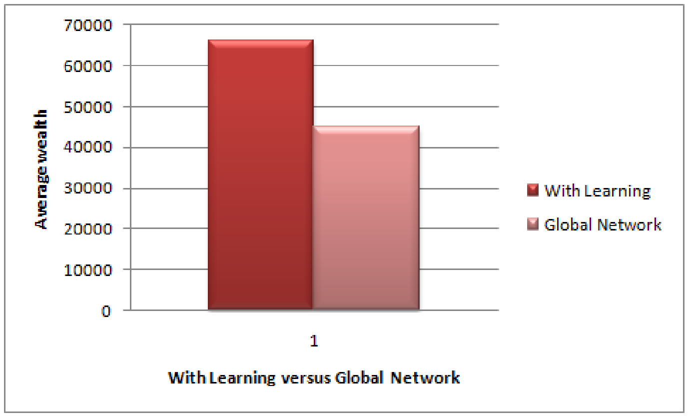

Figure 7.3: Chart illustrating effect of learning on average wealth compared to the Global Network

as a whole becoming better off. The fact that society as a whole becomes better off is explained by returning to a fundamental concept in the model adopted, namely the Pareto trading paradigm. Since every agent will only engage in trade that makes them better off, every trade makes society as a whole better off also.

However, the interesting point is the creation of new wealth. Wealth that didn't exist before all of a sudden does. So this begs the question exactly where has this wealth come from? Perhaps it is purely down to the opening of borders. To conclude this however, we need a point of reference. Learning moves society to a more global position, where more agents are connected to other agents, and in essence borders are broken down. So the most suitable point of reference seems to be the Global Network. If the wealth creation is purely from the ability of agents to search for trade partners in a larger subset of the population, one would expect that the average wealth with learning in the Small World and Local Connected networks would approximate that of the Global Network, ceteris paribus. After conducting this experiment, this was shown to be the case. However, Figure 7.3 illustrates the effect that decreasing the sight of agents in the Global Network to 20 (the same number of agents as in a district), and allowing agents in the Global Network a memory of size 10 (the same as with learning), has on the average wealth. Now, average wealth of agents with learning enabled relative to average wealth in the Global Network is approximately 32% higher, which is also proved to be significant via a means test.

So there is a success of the implementation of learning, manifesting itself in wealth creation when compared with a global network with the same restrictions on sight and memory. This can be explained by the addition of strategy and less search. When learning is enabled, agents are able to locate through their contacts, the best partners. This means that, although you give up time to

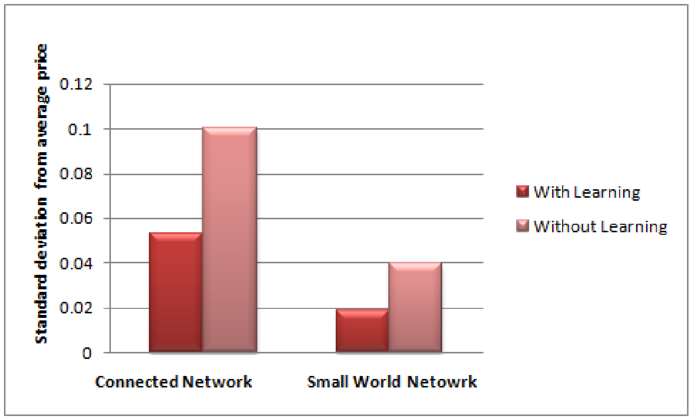

Figure 7.4: Chart illustrating standard deviation from average price, with and without learning

learn, agents in your memory are well suited to you. It also means that if circumstances change, and someone you used to depend on no longer offers good quantities or prices, you can learn of someone else who does. This means, when an agent doesn't decide to learn, and searches through its memory for a trade partner, some of its "memory items" are of agents specifically picked to benefit the agent in him. However, in the Global Network, agents in the memory are simply random encounters of random trades. They may have been one off trades, they may have been sub-optimal. The point is, when selecting randomly, and the agents in your memory are really still a product of random but lucky selections, the probability of finding the best agent for you in the network is low. Hence, you are likely to make less profit from trade partners in your memory. On the other hand, with learning, as time progresses the probability that an agent you learn of is the best agent for you in the network increases since the knowledge in the network increases. This leads to new, previously unwitnessed wealth being created through strategically seeking suitable trade partners as opposed to hoping you happen to find one.

This wealth creation can be seen as a suggestion that the argument avidly debated on the topic of economic globalisation, globalisation makes everyone better off, is in fact potentially correct. However, this is in a far more simple idealised world in which barriers to trade do not exist, subsidies for local produce do not exist and so on. Nonetheless, in this ideological (and perhaps what economic globalisation strives toward) world, it is the case that becoming strategically, or intelligently global does in fact create wealth, benefiting society as a whole on a relative scale. In addition, it is important to note that prices could effect this conclusion. If average wealth increased and prices increased, then the increase in average wealth would be cancelled out by the increase in prices. However, prices were never seen to be higher when learning was employed and thus the conclusion stands in the context of the simulation.

Moving on to price dispersion, the chart in Figure 7.4 illustrates the massive fall by one half in standard deviation from average price when learning is enabled for both the Local Connected and Small World networks. This is also shown by the graph depicting pricer per iteration, where the convergence is easily noticeable. Since optimal trades can be made, goods flow in an efficient manner around the network, removing price dispersion across districts to reach a stable price with small oscillations.

7.3.2 Emerging Globalisation

Architecture

In order to measure the emergence of globalisation, networks depicting districts as nodes and directed edges to correspond to imports and exports were created. If district A imports from district B, then their is an arc from A to B. If district A exports to district B, then there is an arc from B to A. This means that the direction of arcs reflects cash flow, or the flow of Good 2.

Upon initially creating this network, with edges apparent simply if there is trade, most countries engage in trade with most other countries. In order to measure the centralisation of the network, the ratio of edges was taken relative to those edges apparent in a start network of the same size, for both the import and export graphs. The Centralisation Index thus measures the degree of variability in the nodes of the network as a percentage of that as in a star network of the same size.

The chart in Figure 7.5 shows the centrality at the 0% trade threshold (trade exists) and the 5% trade level. At the 0% level, arcs exist from i to j if district i imports from district j. At the 5% threshold, arcs exist from district i to j if at least 5% of your imports come from district j. Centrality is shown for both the import and export networks. The significance of the 5% threshold stems from the fact that there are 20 districts. If every district traded an even amount with every other district, then the amount of imports and exports from district i to district j would be 5%.

From inspecting the chart, it is clear that as the threshold increases, the networks become more centralised. When the threshold is at 0% every virtually every district engages in some form of trade, be it a small or large amount, with every other district, implying an extremely decentralised network. In addition, the index for both the imports are export networks differ insignificantly. However, by increasing the threshold to just 5%, dramatic changes in the structure of the network occur, and the biggest change by far occurs for the import network.

The increase in centralisation illustrates that when the threshold moves to a more meaningful level of trade, all countries export to a small number of partners. The centralisation of the import network however is far more noticeable. When the threshold is set to a meaningful level of trade, the majority of imports have the same destination - one of a small subset of the entire population of districts. These districts take the bulk of imports that are exports for a large

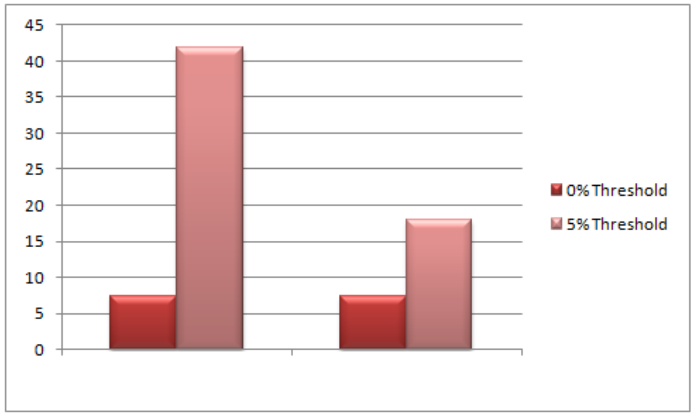

Figure 7.5: Illustration of Centrality as a percentage for the import and export networks using different thresholds

number of countries. In this way, these districts can be seen to act as centers of gravity, attracting goods to their region. This is analogous to the emergence of a core-periphery structure, in which the core are the districts importing the most, and the periphery the others exporting. This in turn means that the core is in a position to exercise a large amount of influence on the districts of the periphery.

With respect to the realism in this finding, it is actually quite astounding. This core group of districts act in a similar way to the G8 countries, buying vast amounts of goods from countries across the rest of the world, the periphery. As such, the core or G8 are able to exert influence on to the periphery - the developing world due to the dependency that these countries have on the core of the network. This finding further enforces the power that the G8 have, and although the simulation only deals with economic consequences of globalisation, in the real world this core-periphery structure can explain for many things from why political bullying works, to the fact that global decisions of world depend on the decisions of the G8 leaders - the core of the network. It is hugely exciting that through the implementation of such a simple method of knowledge diffusion, structures readily apparent in the real world emerge, adding further justification into the insights that agent-based modelling can provide to the field.

However, the centralisation wasn't quite as high in the import network as is witnessed in reality. From the data provided in the paper "The Architecture of Globalisation", [12], the Centralisation Index was computed to be in the region of 77%, so what is witnessed in the simulation is considerably lower, yet still very much apparent, centralisation. This could be explained by the simplicity of the simulation relative to the complexity of globalisation. In the real world, we have trade barriers from taxes, some countries are wary of international trade, especially free trade, and as such it may be more difficult to enter the core of the network. In addition, the most likely biggest difference is the role of governments. Countries already in the core seek to maintain this position, and as such can attempt to prevent countries entering so as they do not lose out on the benefits - they do not lose their power. This notion of those with the power can exert more power and as such keep the power is not apparent in the simulation. Rather, it is a world in which the most efficient districts rule, free trade promotes efficient allocation of resources, and there do not exist barriers to entry.

This notion can be seen in two lights. One in that it is a flaw in the model - a sign that fundamental detail is lacking. However,on the other hand exists a more interesting and compelling argument. The model seeks only to capture globalisation in its economic form. Perhaps the realism lacked is not a flaw in the model per-sé, rather it is an indication of the corruption by mankind of the benefits of globalisation. Perhaps it is the complexities and unfairness introduced by human motive and human interactions that results in a world where countries are suppressed by others, to benefit the economic and social well-being of those in power. This is a fairly abstract argument and as such can benefit from an example.

Consider third world farmers, and a British cocoa distributor. They import an abundance of cocoa from these third world farmers. Britain is a member of the G8 - a member of the core, and the country of the third world farmer, Ghana, is the a member of the periphery. Britain has power to buy cocoa from the farmer for a painfully low price. This does not directly damage the farmer - after all if it wasn't for Britain he wouldn't be exporting this batch of cocoa. However, imagine Britain imports an abundance of raw materials from Ghana. The fact that we pay such a low price prevents Ghana earning much for its produce, and as such does not have the funds to import from countries such as Britain and the rest of the world. It is thus unable to enter the core. One may argue that if Ghana will accept the price, Britain should pay no more, and this seemingly unfair price is just the result of market efficiencies. Although in some respect this is true, it is the fact that Ghana cannot escape its dependence on the core that is the point. It has no choice to accept this offer, as without it, the country is definitely worse off.

So in the real world, members of the periphery being dependent of the core forces a vicious cycle in which it is difficult to escape such dependency due to the influence the core can have on the periphery. In the simulation, free trade means this influence is lacking. Members of the core have no way of suppressing the periphery - the mechanisms don't allow for it. They rise to the top through having earned enough Good 2 to make large purchases. If another district is able to do this, nothing stops it. Rise of another district may take trade from others, but this rise is possible. In reality, some governments go out of their way to protect economic well-being in their own country, at the direct expense of others.

A good example of this would be the American governments method of subsidising local farmers, and applying heavy tariffs to foreign produce. America is not as efficient in farming as some parts of the world, foreign countries offer

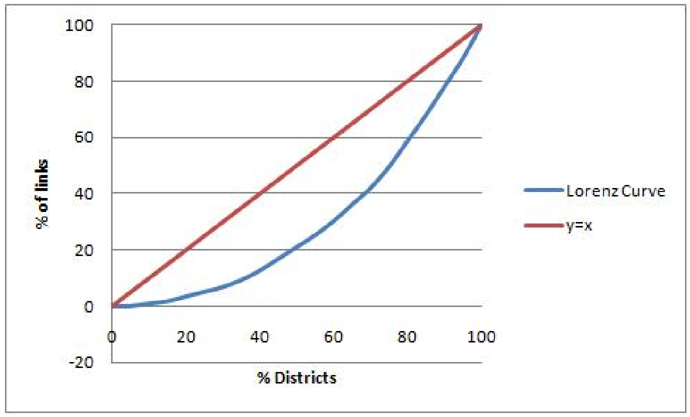

Figure 7.6: Lorenz curve showing the distribution of links in a network

far better prices. The American government knows that few businesses will pay more solely for the reason that it is grown in the United States, and thus they simultaneously make foreign produce more expensive through tariffs, and local produce cheaper through subsidies. This is an example of entirely irrational economic behaviour in what we believe to be an increasingly globalised world. This is in fact illustrative of the fact that we are far from truly global economically. Countries are still concerned about having dependencies on countries for items such as food, and governments still favour the short term interests of the national, not global population. However, the advantages of free trade as it is in the simulation, have been hypothesised in economic theory from as early as 1871, where David Ricardo wrote a book titled “On the Principles of Political Economy and Taxation” which bore the idea of comparative advantage. Put simply, comparative advantage addresses the fact that free trade means that if one nation is better in production of a good at another - it can do so with lower cost, offering lower prices, then they should focuses their production on this good. This enables efficiencies and prices that couldn’t be realised if the country less efficient in production were to produce the good. America and its farming is an example of a direct contradiction of this idea.

In addition to analysing network structure based on centrality, the Lorenz curve was employed to enable the graphical visualisation of distribution of edges across the network of districts. Lorenz curves were computed for both import and export networks and can be seen in Figures 7.6 and 7.7 respectively. These again illustrate the differences in the network structure for imports and exports.

The graphs can be interpreted in a similar fashion to those of wealth distribution. Along the x-axis we have the percentage of districts, and along the y-axis we have the cumulative percentage of total links apparent in the network, stemming from these districts.

The graphs further highlight the more uneven distribution of edges in an import network - with a few nodes holding the majority of the edges, juxtaposed

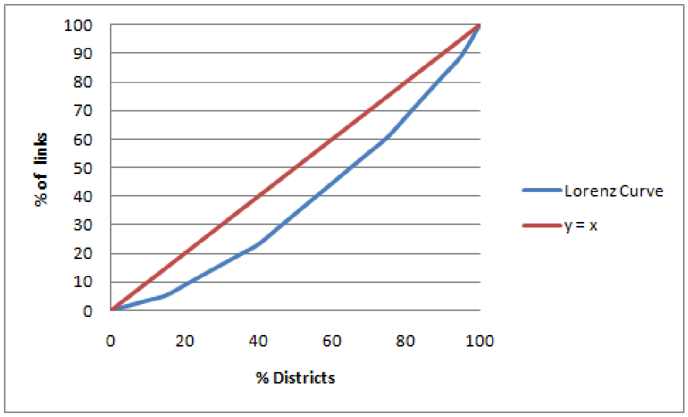

Figure 7.7: Lorenz curve showing the distribution of links in a network

with the export network which is far closer to an equal distribution. This is a perhaps simpler graphical way to understand the effect of a more centralised network on import and export networks. Again the distribution is far more fair than that of the real world, for reasons already discussed.

Correlation between wealth and deficits

In economics, on of the measures of economic integration is the looking at a countries trade balance. The trade balance of a country is simply its exports minus its imports. If the trade balance is positive, it is known as a trade surplus - you export more than you import; if it is negative, it is known as a trade deficit - your import more than you export. Traditionally, in economics a trade deficit is thought of as a bad thing, although this is not set in stone. In the simulation however, the opposite was consistently true. A trade deficit always led to higher wealth. Instead of a trade deficit being an indication of high debt, it is an indication of power. The reasons for this trend in the simulation is two fold. Firstly, debt doesn't exist and as such, a trade deficit says nothing about the debt of a country. Instead it reflects that the country is in a position spend money to buy more goods and create more wealth. As previously discussed, the greater wealth cannot be attributed to the prices differences, since these are negligible across districts. However, although the simulation does not have debt, it does not imply that this observation is unrealistic. On the contrary it turns out that the wealthiest and most powerful countries of the world have the largest trade deficits. Judging by what we have witnessed in the import networks, this is hardly surprising. The core of the network are destination for the majority of imports, and the core of the network in the real world is the G8. The G8 import the largest quantity of goods, and thus have the largest trade deficits. The reason for the high level of imports is out of the scope of this project, however, it is likely that the trade deficit is not just an indication of the a countries level of debt, GDP growth or unemployment, but also a good estimator of its position and role in the international trade network.

Creative Destruction

We have witnessed that districts rise to a position in the core through the diffusion of knowledge of agents offering the best deals. What has been overlooked however, are these new bridges, and the importance of these new cross district traders in the emergence of globalisation. In fact, inspecting this gives some fairly exciting results. For clarity, to demonstrate the results, one simulation will be looked at.

Creative destruction is a notion in economics to do with innovation. It describes the process by which innovation can result in entrepreneurs entering the market place and fueling economic growth, despite it destroying value of other established companies. In the context of the simulation, this can be seen as an agent losing their position in a market place to a more proficient agent. Although the agent who suffers from this, and his wealth may well be reduced, it is this process that makes society as a whole better off. An agent should not be given a valuable position in a market place if they are inefficient. Recall one of the things learning sought to eradicate was the unfounded fortune of crossover agents. They, for no reason, enjoyed a position in the market, that firstly they may not be able to fully exploit, but secondly, do exploit despite there being a non-crossover agent who could do it better.

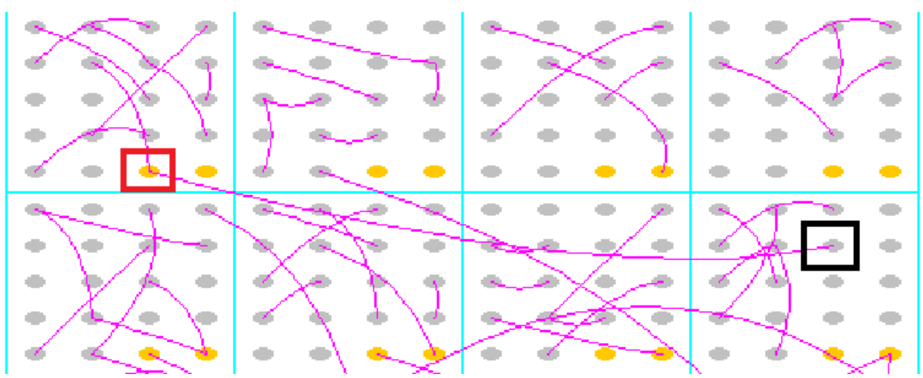

Figure 7.8: A crossover agent and his trade partner

In order to illustrate this pictorally, a graphic from the simulation was extracted. It may look like there are few trades, however, the image is of only one iteration. This is to make the arcs more clear however also means the extent to which creative destruction occurs is lacking. In the simulation, the trade occurring between the crossover agent and the foreign agent occurs more than once.[2]. The image provided in Figure 7.8 shows a crossover agent (highlighted in red)

[2]A temporary output to CSV was implemented purely for the purpose of looking for creative destruction, and this was studied to decide which simulation to use as an example.

and its trade partner in black, near to the beginning of the simulation - at 50 iterations. Notice that this crossover agent has an arc entering it, representing a trade (not necessarily initiated by him). Other agents in the district do not trade much with agents outside the district, we are close to the beginning of the simulation and knowledge has not been propagated across the network.

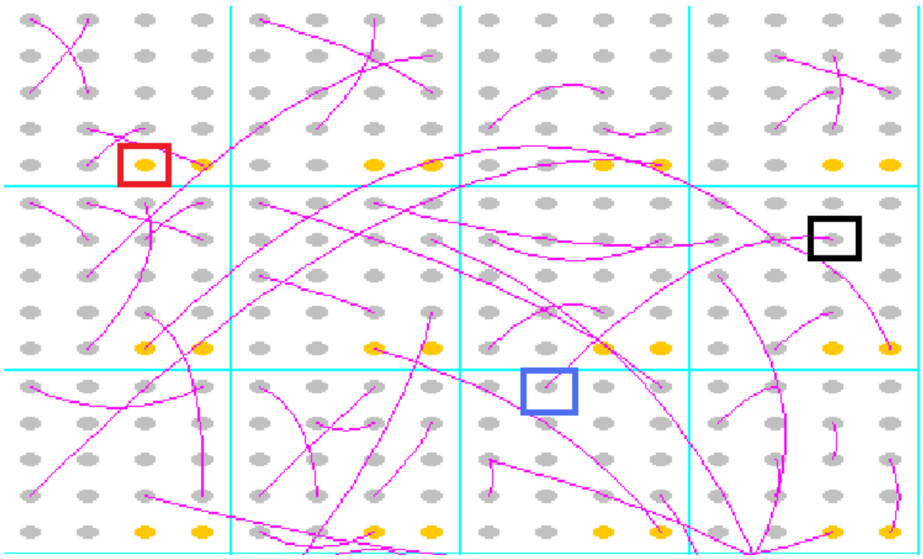

Figure 7.9: The fall of the crossover agent and the new trade partner

Figure 7.9 shows the same simulation, only this time 500 iterations later. By this time, a fair amount of knowledge has diffused through the network and agents are more strategic about with whom they engage in trade. Now three agents are highlighted, the crossover agent again in red, and in addition a non-crossover agent in blue and the trade partner in black. Notice that the node that had an arc from itself to the crossover agent (highlighted in black) now have none to the crossover agent, and instead one to the new agent. This new agent has emerged as a more efficient and effective trade partner, with better prices and more to offer. Also notice that this new district bridge is not a crossover agent. In fact, for the duration of the simulation, the once powerful crossover did not engage in cross district trade again.

Not only does this illustrate that learning has provided new efficiencies to be witnessed, but in addition is an indication that creative destruction is not a notion that requires innovation necessarily. In the simulation, there is no concept of innovation, yet this is still witnessed. Rather, creative destruction can also be realised through free markets and free trade. They allow for those who are most proficient in a job to rise to that job. As such, not only do they become more wealthy, but the wealth of society as a whole is increased due to better prices and quantities being available to other agents.

7.4 Conclusion

Learning has proved to be successful in facilitating the proliferation of knowledge through the network. In addition, it has allowed for globalisation to emerge in a more decentralised manner than the real world, yet still with a clear core-periphery distinction. Questions have been raised about the trade barriers in the real world perhaps preventing the benefits of globalisation to in fact be felt globally. Unfortunately, further investigation into answers to these questions are out of the scope of the project.

Creative destruction has been witnessed as agents who deserve a position facilitating global trade have earned them through reputation and reliability, and it is suggested that the notion of creative destruction need not be restricted to the requirement of innovation being the causal factor. For further justification to this argument, we can revisit the American farmers. If the American government were to engage in "laissez-faire" economics - leaving the economy to market mechanisms by refraining from government intervention - then it is virtually definite that a large portion of the American farmers would be bankrupted, and better, cheaper providers of the same good could take their place. However, this brings us to the limitations of mankind to the progression of globalisation. Free trade requires trust, and it requires governments to act in the best interest of the world, not solely their nation. As such, unlike in the simulation, some of the efficiencies, fairness and advantages of globalisation cannot be realised for everybody. Flaws in ourselves hamper the realisation of benefits in becoming truly global.

Chapter 8

Evolution as a method of gaining insight

In the simulation, evolution can be employed in order to gain insight in to the model. It allows a user to see what makes a fit agent - what character(s) does the genetic algorithm cause agents to evolveinto given the initial conditions of the simulation? This provides insight into what effect initial conditions have on the evolution of the simulation, and under which particular conditions certain strategies are optimal for agents. It is important to note that it is not reflective of biological evolution. Instead, the genetic algorithm creates a new population from the current population, and the current population exits the simulation. This process is repeated to evolve an increasingly fit population, and each new population is examined with respect to the others to find general trends in the evolution of agents.

Let us begin by explaining in detail how the genetic algorithm works, specifically with respect to the simulation.

8.1 Genetic Algorithms

Genetic algorithms encode a potential solution to a particular problem, and the method of attaining this solution is analogous to that of evolution. In the Background section the reason that evolution is such a good method for complex optimisation problems in large design spaces was explained.

The genetic algorithm begins with a population of randomly generated agents. These agents are allowed to produce, trade and so on for a period of time. Each agent is then assigned a value based on its fitness for solving the problem; in this case, problem solving ability is gauged according to how high their utility is and how wealthy they are. It may seem that utility and wealth are two ways of displaying the same thing. However, since prices can vary, and are not necessarily 1, measuring wealth also captures value. This means an agent with a relatively large stock of goods who has a lot of Good 1 can be considerably wealthy if the price of Good 1 is high, i.e. if Good 1 is in high demand. Therefore it is necessary to consider both sides - utility and the value of their assets. Upon each agent being assigned a fitness value, agents are selected for breeding. The probability of a fit agent being selected is higher than that of a worse off agent, but the worse off agent still has a chance of being selected to breed. Each agent that successfully makes it through (an agent can make it through multiple times), is randomly chosen a partner with which they will be bred. This pair of agents, to whom we will henceforth refer as parents, are then bred to generate two new offspring using a process of crossover and mutation. These offspring make up two agents of the new population. Once the new population contains as many agents as the first population, it is initialised and the first round of the genetic algorithm is complete. These agents then produce and trade and so on until they are evolved into the next population. The process continues for a finite amount of time specified by the user.

8.2 Implementation

Firstly, we have to define what changeable aspect of an agent should be considered. This will determine what is bred between the 2 agents. The following is what is evolved:

- Amount they can produce of Good 1
- Amount they can produce of Good 2
- The size of their memory
- How many agents they can search through to find a trade partner

And the fitness of agents is determined by:

- Wealth
- Utility

The new population will be the same size as the initial previous population. Next, the implementation of the genetic algorithm will be explained and this will cover the stages below:

1. Assign a fitness score to each agent
2. Pick, based on the fitness scores, the agents that make up the parents of the new population
3. Generate the offspring for each set of agents
4. Initialise the new population, and run restart the simulation

8.2.1 Fitness Scores

As in the case of memory and learning, the multi-attribute problem leads us to generate values for each of the two variables, wealth and utility, and combining the weighted values to attain a fitness score. The value for wealth and utility are calculated as follows. A value closer to one is used for high utility and high wealth:

- Value for utility for agent i[1]:

  $$U^i = \frac{v_i}{\mu}$$

  where

  $$v_i = \frac{u_i - u_m}{u_r} \qquad \mu = \frac{\sum_{i=1}^{n} v_i}{n}$$

  and
  u_i is the utility of agent i, u_m is the minimum utility of the population, u_r is the range in utilities, and n is the number of agents in the population.

- Value for wealth for agent i, W^i is the same for that of utility, but substituting utility for wealth.

The aggregate value for an agent is given as:

$$V^i = 0.6 \times U^i + 0.4 \times W^i$$

Then, each agent is assigned a new value, as a proportion of the total values of all the agents. Formally:

$$F^i = \frac{V^i}{\sum_{i=1}^{n} V^i}$$

where
n is the total number of agents.

F^i therefore describes what proportion of the total fitness of all the agents agent i has, which is used in selecting the agents to breed.

8.2.2 Agents to make up the new population

Now each agent has been evaluated in terms of its fitness relative to other agents, and has been given the proportion of the total fitness that his fitness represents, it is possiblez to generate a new population. As previously mentioned, even the less fit agents have the opportunity to breed. The selection process is achieved through employing a technique known as *roulette wheel selection*. Roulette wheel selection works by giving fitter agents a higher probability of being selected, and

[1]In the canonical genetic algorithm the fitness score is the evaluation of the agents fitness over the average evaluation[15]. For a detailed tutorial on genetic algorithms see [15] in the bibliography.

poorer agents a lower probability. It works as follows. Imagine the total of the values,

$$V = \sum_{i=1}^{n} V^i$$

to be a pie, and each agent's proportion of total fitness,

$$F^i = \frac{V^i}{V}$$

to correspond to a share of this pie. If an agent has $F^i = 0.1$, this indicates that he is entitled to 10% of the whole pie. Each agent now has a slice of the pie, and the size of the slice corresponds to their value of F^i. Now, we rearrange the agents in terms of the size of their slice, in ascending order. Each agent is also now ranked by its index in this new ordered list from 0 to n. We now have a pie that looks something like Figure 8.1.

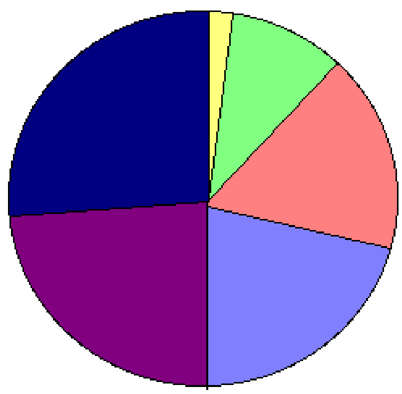

Figure 8.1: Illustration of a pie corresponding to a set of 6 agents

Selection is carried out by generating a random number, and the agent whose range captures this number is selected to go through. The range that each agent has depends on its score and its rank. Formally, the range, L^i to T^i belonging to agent i with rank r and score F^i is given by:

$$L^i = \sum_{j=0}^{r} F^j \quad T^i = L^i + F^i$$

Random numbers are generated, and the corresponding agent is selected until the number of agents selected is equal to the number of agents in the new population. Each agent in the list of "successful" agents is now randomly allocated another agent to be its partner. Once every agent has a partner, offspring are generated.

DNA P1	1	1	0	1	0	0	1	0	1
DNA P2	0	0	0	1	1	1	0	1	0
r	0.6	0.4	0.1	0.2	0.8	0.9	0.7	0.3	0.6
DNA C1	0	1	0	1	1	1	0	0	0
DNA C2	1	0	0	1	0	0	1	1	1

Figure 8.2: An Illustration of crossover between two agents *P1* and *P2* generating children *C1* and *C2* based on random number r

8.2.3 Generating offspring: Crossover and Mutation

The next step of the algorithm is to generate the new population of agents. As a reminder, the genes that characterise an agent are:

- Amount they can produce of Good 1
- Amount they can produce of Good 2
- The size of their memory
- How many agents they can search through to find a trade partner

Each of these is an integer, and thus can be represented as a binary string. Concatenating these strings gives us the "DNA" of the agent, as shown in Figure FIG HERE. Crossover is the act of combining the DNA of both parents to generate offspring. Here I use a method known as uniform crossover. Here, each of the two children have an equal probability of inheriting a single bit from a particular parent. Let P_1 and P_2 be the two parents, and C_1 and C_2 be the children of those parents. For each bit in the DNA of the parents, generate a random number r^i.

- If $r \leq 0.5$, then the *ith* bit of P_1 corresponds to the *ith* bit of C_1 and the *ith* bit of P_2 corresponds to the *ith* bit of C_2
- Otherwise, the *ith* bit of P_1 corresponds to the *ith* bit of C_2 and the *ith* bit of P_2 corresponds to the *ith* bit of C_1

Figure 8.2 gives a graphical example of crossover between two agents. In reality, not all combinations are valid. For instance, if the upper bound on how much an agent can produce is 30, then 5 bits are needed to encode all possible values. However, this gives 31 possible values. Therefore checks have to be done to ensure the value realised is with in the bounds. In addition, if the upper bound is 30, but the agent can only produce 10, in binary this would be *1010* which is only 4 characters long. In this instance, this would be padded to *01010*.

The next step is mutation. The mutation rate is the probability that a particular bit will be flipped, and in the simulation is set to 0.07. For each bit of the children, generate a random number uniformly in the interval [0,1). If this is below the mutation rate, flip the bit and move to the next bit, otherwise simply move to the next bit and repeat the process. The bit string is then split up back into the 4 attributes that characterise the agent.

Having generated the genes for the new population, they are intialised. The number of crossover agents in the new population will be equal to the number in the initial population and which agents become crossover agents is decided randomly.

8.3 Evaluation

Simulations were conducted and agents evolved in order to gain insight into the model, particularly to determine whether restricting the production of goods, learning and memory increased the value of trade or simply made it necessary or more available. This can be achieved through analysing the level of trade as well as strategy shifts displayed in trends of specialisation of agents. In addition, changes in both wealth and its distribution are examined.

Without any of the extensions made to Wilhite's original model, the genetic algorithm caused agents to become pure producers. In addition, trade fell to 0.8%. This is illustrative of how much better production intrinsically is in this model. Agents do not need to change, and the lack of necessity for trade in the population means agents simply become proficient at production and the production functions of agents tend to their maximum. This is due to the fact that there is nothing but production functions to vary and, as such, little scope for any further direction of evolution.

Memory and learning, however, offered more interesting turns of events. The evolution of agents is very similar in both circumstances, but differences in the specialisation of agents is apparent.

Let us begin with the similarities, namely wealth distribution and average wealth of society. For both learning enabled and solely memory enabled, the wealth distribution in society becomes incredibly even. Figure 8.3 illustrates this, as the Gini Coefficient falls sharply in the first four iterations of the evolutionary algorithm. As agents evolve, production functions are optimised. Agents tend towards having very similar abilities, particularly in production, despite the mutation employed in the algorithm. This means that the distribution of wealth practically reaches perfect equality - convergence happens too readily.

In addition, the average wealth of society increases; indeed, over the course of the algorithm it is actually doubled, as illustrated in 8.4. Again, this can be contributed to the proficient producers that develop. The possibility of it being attributed to efficiencies in trade from the employment of memory and learning can be dismissed upon inspecting the graph describing the level of trade when learning is enabled. The amount of trade falls rapidly within the first 4 iterations of the algorithm before levelling off, shown in Figure 8.5. With a

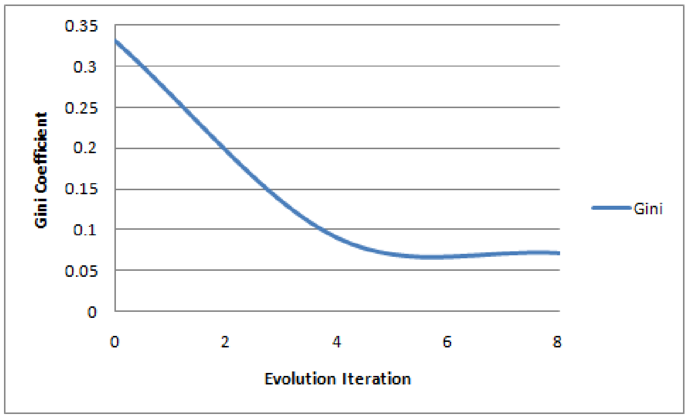

Figure 8.3: Change in distribution of wealth, measured by the Gini Coefficient, as the genetic algorithm progresses

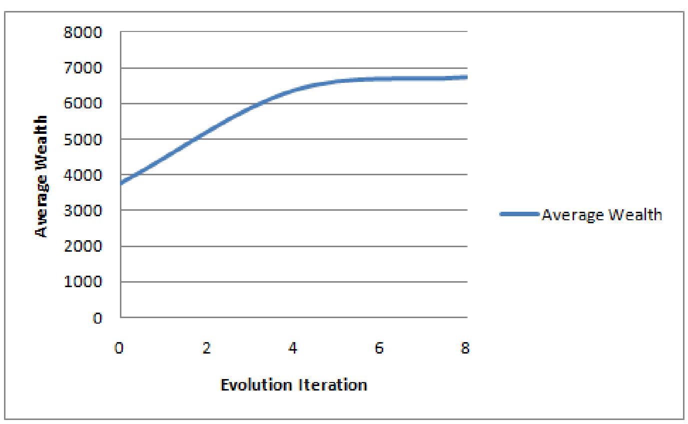

Figure 8.4: Change in average wealth as the genetic algorithm progresses

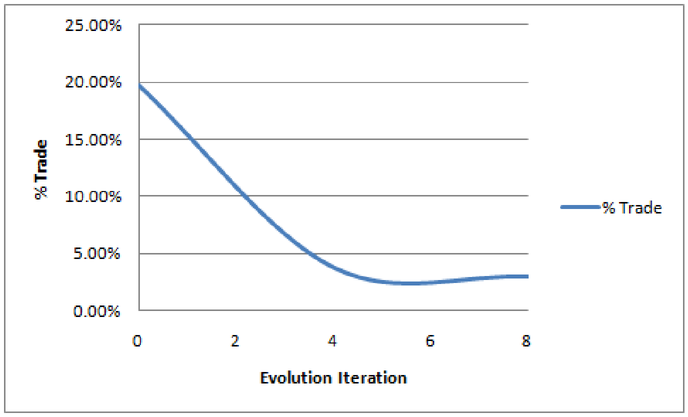

Figure 8.5: Change level of trade as the genetic algorithm progresses

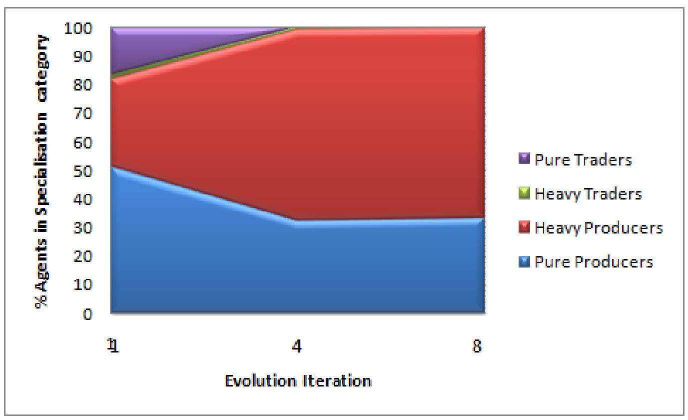

Figure 8.6: Change in agent specialisation as the genetic algorithm progresses

world of proficient producers, trade simply is not necessary. There is less value in trade than there is in production, and as trading has always been a last resort, the algorithm works to remove this flaw which acts only to restrict equality in the economy.

This shows that as the evolutionary algorithm progresses, since the traders are the poorest agents, often by quite a long way, the probability of them getting through to the next round begins to fall. However, this does not address the question of where they actually grow.

This can be seen by examining the specialisation of agents. Figure 8.6 shows the specialisation of agents with the progression of the evolutionary algorithm. As agents evolve, the pure traders and heavy traders disappear. They are no longer worthy of existing in the economy. However, like with memory or learning, this time it is the heavy producers that triumph. All agents evolve

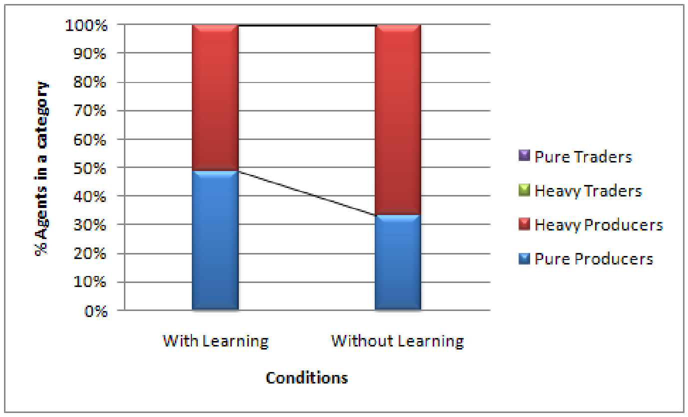

Figure 8.7: Comparison of change in agent specialisation evolving with and without learning

to favour production, but the percentage of pure producers actually declines, as heavy producers replace both pure and heavy traders, and pure producers. Comparing the breakdown in specialisation at the end of the simulation with learning to the those with only memory, a fairly large difference is apparent. This can be seen by inspecting Figure 8.7. In both contexts, by the time the genetic algorithm is finished, there are no pure traders and a negligible number of heavy traders. However, learning has proved its success in facilitating trade networks. Despite there being only a small increase in the evolved level of trade when learning is employed, is significant (via a means test) when compared with the level of trade with just memory. From Figure 8.7 you can see that with learning, there were fewer pure producers in the end. This is indicative of the fact that learning offers beneficial trading opportunities. This in turn means that although agents are evolving to become producers, learning offers more opportunities for trade through the ability to both remember and seek out a trade partner. It seems that learning is the extension that makes the most difference to the value and potential benefit of trade.

8.4 Conclusion

In conclusion, the genetic algorithm validates the idea that production is unequivocally favoured, and that the evolution of the simulation is largely based on the ability of agents to produce goods. Nonetheless, it seems the ability of agents to learn permits some value to be added to trade. Agents production functions do not necessarily evolve to be 30, as some agents can rely on the occasional trade. The fact that production functions are higher also means that trades that take place can be with bigger quantities and, with learning, with a bigger quantity. This in turn means that fewer trades have to occur in order to even out stockpiles and maximise utility.

The results also highlight the need to create value of trade in the simulation, and this is to be addressed in the following section.

Chapter 9

Implementation

The application is a web application implemented using a combination of Java and Python for the back end and HTML, Javascript and CSS for the front end. The application is split up into three main sections: *simulation, user interface, and analysis engine*. The user interface is the users access point and allows creation, running and downloading results of a simulation. The analysis engine is responsible for data analysis, producing graphs, charts, tables and PDF generation. Both of these sections will be discussed in detail in the following section, however for now our focus lies on the implementation of the simulation.

9.1 An Alternative Design Choice

One of the aims of the project undertaken was the formulation of a piece of software that offers a novel approach to the design and implementation aspects of constructing an economic multi agent simulation. For this reason, although research was carried out on alternative platforms that could be leveraged, it was decided not to use these since it would remove the opportunity for experimentation with the implementation.

However, in this section the most appealing platform studied will be outlined, namely JADE. This is a well established Java based platform and has extremely impressive features as well as a diverse range of applications.

9.1.1 JADE

JADE is a platform developed entirely in Java that enables the construction of multi agent system. Each agent is run on its own thread, concurrency, parallelism and cooperative task scheduling being handled by JADE. Agents can be physically distributed on hosts and on each host there is only one Java application being executed. Agents communicate with one another via message passing with Foundation for Intelligent Physical Agents Agent Communication Language (FIPA ACL) being the language of choice. FIPA ACL also allows the inclusion of user-defined message parameters where the semantics are not defined by FIPA enhancing extensibility and ease of customisation. [31]

In addition to taking care of the majority of distribution and threading issues, JADE also offers an extensive GUI. It allows the user to manipulate the system at run time through the ability to start, restart and stop agents. In addition agents can be tracked through message sniffing during the simulation allowing the exposure of the internal state of subsets of the system. It also offers assistance in debugging techniques which are often complicated when dealing with distributed systems. An example would be the Dummy Agent which allows inspection of message exchange between agents. It enables validation of an agent description prior to integration in to the MAS and in the event that an agent is malfunctioning the Dummy Agent facilitates interrogative testing. In addition it is possible as the user to construct, send and receive messages from agents in the simulation. [30]

However, JADE's focus is primarily on the physical distribution and computational autonomy of agents. Although this is appropriate for the implementation of MAS, in the model I am extending, the architecture is less appropriate. This is due to the fact that in my model, agents get to pick an action sequentially, and rounds are repeated iteratively. Hence the concept of message passing and real autonomy is somewhat redundant with respect to the synchronous single threaded architecture I planned on undertaking. It may be viewed by some that you are not creating a truly autonomous MAS without having physical distribution of computation and true agent autonomy. However, I believe, having done research in to this debate that it is purely an implementation issue and with the timescale at hand, and the model being implemented, it is reasonable to adopt a method of pseudo autonomy with no physical distribution. I believe the implementation choice will have little to no impact on the result of the simulation since the logic by which agents make decisions are homogeneous between the two options.

9.2 The simulation: Java & Python

Agent Based Modelling generally requires the use of Object Oriented programming in order to be able to model agents and their environments as entities. Hence, the two languages that the simulation is to be written in are both object oriented.

9.2.1 Java

Java is to be used to model the artificial world and will be primarily for data storage. The Java section of the application defines what an agent is, what a world is and what a district is. The idea of the Java code is that it is the core of the simulation - a description of what it is and not the way in which agents interact with each other. The reasoning behind this is that these are attributes of the system that are not to be experimented with, and should remain unchanged across simulations. The logic contained in the Java code is purely “safety logic” in that it facilitates functionality with no impact on the result of the simulation. However, the project is somewhat experimental. There is a need to try different ways of agents interacting. For instance, in the initial model, agents have a choice of only two interactions, production and trade. As this is extended and more interactions are added, it is desirable not only to be able to draw comparison between the two, but also to avoid the loss of the other more simple set-ups I implemented. It therefore seems that there is need for a dynamic scripting language, namely Python.

9.2.2 Python

Python is an object-oriented, dynamically-typed interpreted language. It allows for great flexibility as well as the bonuses of object-oriented design. It is also easily integrated with Java projects through the use of Jython[1]. In the projects objectives it was noted that one aim of the project was to allow for further work to be carried out. Jython allows just this through enabling embedded scripting, i.e. end users are able to write scripts to enhance functionality to the application.[28] Not only is this useful for users or developers in the future, but it will also be so during the development of this specific project. The use of Python will allow to turn extensions “on” and “off” as previously mentioned through the use of loading different scripts or not loading various scripts describing the logic of the simulation. Another incentive to use Jython is that time is of the essence for this project and for it to be useful to other users it should be as efficient as possible. It is therefore imperative to work with a language that increases developer productivity. Python programmers are typically 2 - 10 times faster than Java programmers. Scripts will be used for all interactions of agents:

- Trade: Search, and Negotiation and Exchange will be separate
- Production
- Initial endowments

This allows for the experimentation with regard to how agents make their decisions and select partners etc.

The scripts prevent loss of more simple or different simulations and allow far easier experimentation as well as simply providing organizational enhancements. They make it simpler to run different simulations with different configurations of logic in order to facilitate comparisons across observed macro behaviour relevant to simulations of different scope or logic.

[1] Jython is an implementation of the high-level, dynamic, object-oriented language Python seamlessly integrated with the Java platform.

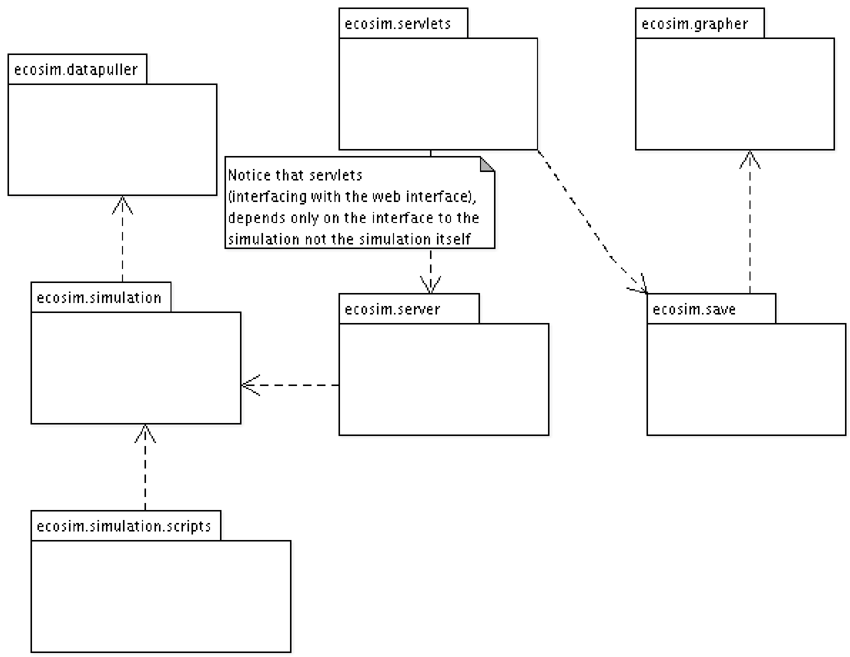

Figure 9.1: Overview of package structure

9.2.3 Architecture

The application was designed with both extensibility and configurability in mind. The class structure of the simulation only can be seen in Figure 9.2, and the package structure of the entire application is illustrated in Figure 9.1.

Initial Model

The Simulation class is responsible for running iterations of the simulation and creating objects to store data. It also acts as an interface to the analysis engine by retrieving information that it may need (such as configuration options). The simulation contains a world which is made up of districts and also houses the script engine. The world is subtly different from the simulation. One way of viewing this difference is that the simulation can be considered to run the world. As has been indicated, the world contains districts, and districts contain agents. In the simulations, Agents are simply actors in the economy. They know how much they can produce of each good, they know where they live, and they know how many people they can search through to find a trade partner.

It can be seen that Agent is a class that has solely a location, and an identifier. A specific subclass of Agent is Worker - this defines an agent who can be active in the economy - they also have a stock of goods and an amount they can produce of each good. If someone were to extend the model to account for

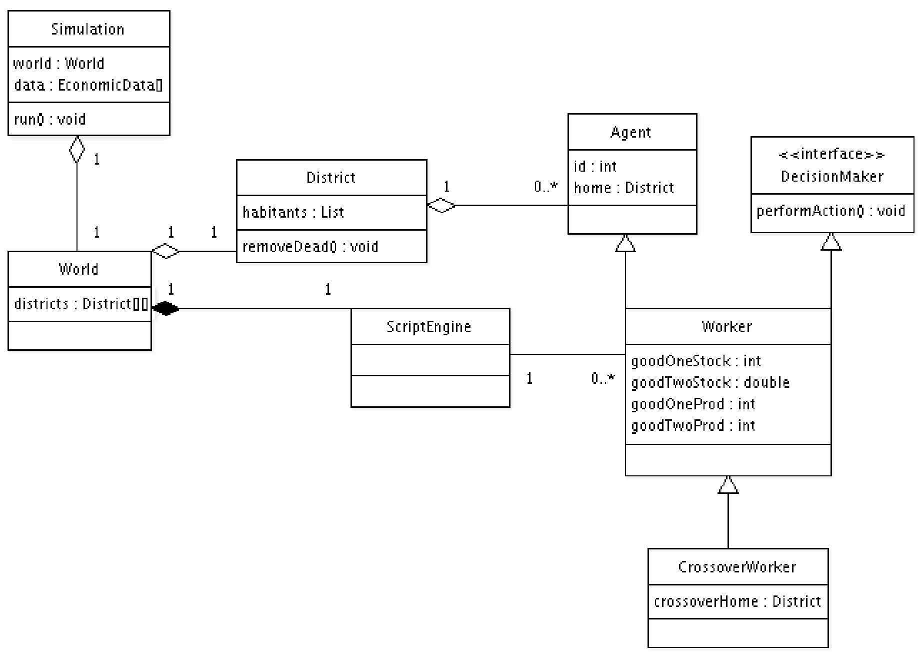

Figure 9.2: Reduced Class Diagram for the Simulation (package ecosim.simulation)

retired agents or any agent inactive in the economy, but existing in the world, this could be achieved simply by subclassing agent. In addition, there is an abstraction between an Agent and a Decision Maker illustrated by the fact that Worker implements Decision Maker whereas Agent does not. This could allow for extensibility again in the sort of agents existing in the world. For instance, if Children were introduced and Parents supported Children, then Children may not be required to actually do anything for a given period and the architecture would easily allow for this.

In the class diagram the Script Engine is evident. This is the only place in the application which actually deals with Jython - the interoperability layer between Java and Python. This encapsulates the scripts, and also holds the responsibility of loading the correct scripts on the basis of the simulation configurations, as well as the packing and unpacking of arguments for interoperability. This allows for the details of configuration and interoperability to be self-contained and thus hidden from the rest of the application, which only interfaces with the Script Engine, meaning it can call the same methods irrespective of the scripts loaded for a particular interaction.

Consumption

Consumption was implemented quite simply by a supplementary Python script which accounted for the fact that an agent, when deciding to produce or trade, also had to consume a stock of goods. The quantity to consume is kept in the script as a global variable in the interpreter and is set prior to running the simulation. The script engine on construction is given a Boolean determining whether or not consumption is "on" in this simulation and loads the Python script accordingly. Again, when an agent has to perform an action, it calls the same method irrespective of whether or not they are consuming.

The only additional implementation in the simulation was the fact that agents can now become inactive in the economy - suffer an economic death, and that consumption should be recorded. Thus the script returns an object containing whether or not the agent dies, and if he did die he adds himself to the list of economic deaths of its district. The district removes the agent from its inhabitants and the simulation records the agents death in the appropriate data store. The script also stores the quantity that the agent managed to consume on a given iteration.

Memory

Memory allowed agents to choose whether or not to remember a certain exchange with a particular agent. Upon performing an action, if the action was exchange, the agent had an opportunity to remember its occurrence. In addition, learning employed the use of memory. When an agent learned of another agent, they had the opportunity to store this agent in their memory. The decision to learn was implemented using a Python script.

If the contact exists in the agents memory, the memory event would be

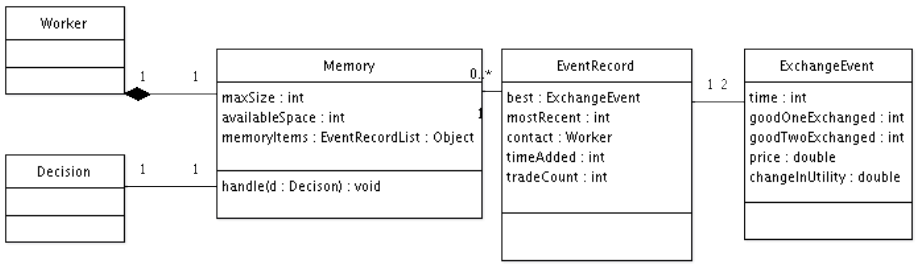

Figure 9.3: Reduced Class Diagram for the Memory Extension (package ecosim.simulation)

updated. Otherwise, the agent had the option of storing the event as a new event (a new agent) in its memory. A Python script implements the decision to store an agent. It takes care of whether to store them, and if the memory is full, whom replace with the new agent. All it creates, however, is a Decision object - it in no way actually performs addition, removal or updates on the memory. This is encapsulated in the Memory object of the agent.

The addition of memory required a “sub-architecture” for storing a memory item. This is illustrated in Figure 9.3. Every agent has a memory, and each memory has a maximum number of events it can store. It also has a list of remembered events - EventRecords. An EventRecord models a memory about a single agent. It holds a reference to who the agent is, when they were added and how many times it has been used. These are used in calculating the aggregate benefit value of a specific memory item when determining who to store, and if necessary who to swap. It also stores two objects called ExchangeEvents. An ExchangeEvent is representative of the occurrence of exchange. It records the change in utility it generated, the time of occurrence, and the quantity of goods exchanged as well as the price. As stated, two exchange events are stored - one for the most recent exchange with the given contact and the other for the best exchange with the given contact.

As mentioned, the addition, removal and updating of events from an agent’s memory is the responsibility of the memory itself. This encapsulates the logic so that the agent can just give a decision to its memory, and the memory handles it itself.

Evolution

The algorithm for evolution was previously explained in detail. As a recap, the simulation runs for a fixed number of iterations. Then each agent is given a fitness value based on its utility relative to the population. The fittest agents have a higher probability of being selected than the less fit agents. For a population of size n, n agents are selected from the current population (these need not be distinct). These agents, let us call them successful agents, are then randomly put in to pairs. These pairs are Parents. Each pair is bred to generate 2 offspring

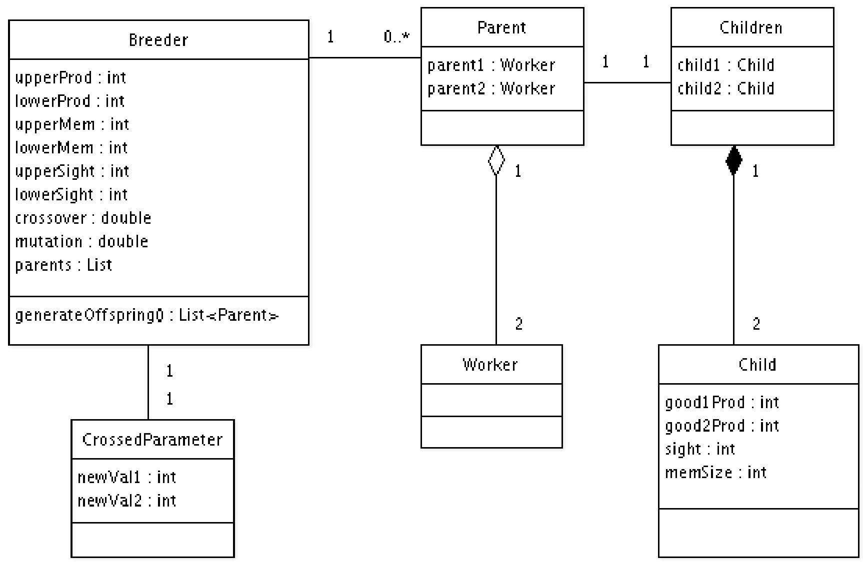

Figure 9.4: Reduced Class Diagram for the Evolution Extension (package ecosim.simulation)

using a technique called uniform crossover (and also performing mutation), and the offspring of all the Parents constitute the new population.

Evolving agents is initiated by the Simulation class since it is aware of the current iteration. However, the repopulation and "resetting" of the simulation - including replacing inhabitants of districts, and assigning crossover agents a crossover district, is delegated by the World class.

The algorithm for generating the pairs of agents to be parents is implemented as a Python script. The crossover and breeding of agents again introduces a new Java "sub-architecture" which is illustrated in Figure 9.4. The Python script passes back a list of Parent objects. The world takes this list and instantiates a new Breeder object. The breeder objects has the list of parents but also has the upper and lower bounds for production, sight and memory size in order to ensure correctness in the crossover process.

The Breeder object generates a list of Children that it sends to the world. A Children object stores the two offspring Child objects of a pair of agents in a Parent. A Child is not an actual agent at this point, it is just the parameters to set up the agent. When the world receives the list of children it is transformed in to a list of workers and crossover workers. The number of crossover workers is the same as at the beginning of the simulation and an agent is randomly chosen to be a crossover worker.

9.3 PDF Generation

PDF documents were chosen as a method of output as it is an effective way of generating a neat report of a simulation. The idea was for it to contain enough analysis in order to perform a proper evaluation of a simulation. The PDF documents contain:

- All information used to configure the simulation
- Graphs illustrating the amount of production and trade over the simulation as well as this as an actual percentage
- Graphical illustrations of trades occurring at quarter points of the simulation (an illustration is given in Figure 9.5)
- Information on wealth and distribution, specifically:
  - Formula for calculating wealth
  - Average global wealth at the end of the simulation
  - The Gini Coefficient at the start, middle and end of the simulation
  - The Lorenz Curve at the end of the simulation
  - A graph illustrating global wealth over time
  - A graph illustrating wealth per district over time
- Information on utility including a graph of global utility over time and utility distribution at the end of the simulation
- A table giving average prices on a district and global level together with standard deviation, and a graph of prices over time
- Information on imports and exports of districts, specifically:
  - A table showing the value and quantity imported and exported for each district together with their trade balance
  - A table showing the percentage of district *i's* exports that go to each district, for *i in 1.. #districts*
  - A table showing the percentage of district *i's* imports that come from a certain district, for *i in 1.. #districts*
- A Pie chart illustrating the percentage of agents falling in to a specialization category
- Information on individual agents, specifically for:
  - The wealthiest and poorest agents
  - The agent with the highest utility belonging to each specialization category

– The poorest and wealthiest crossover agents

Each of these has a table containing data as well as graphs illustrating the agents movement of goods over time, and consumption of goods if applicable.

- A graph showing deaths over time, again if applicable

PDF generation was performed using the library iText. A document is instantiated, and to it Chapters are added. Sections are added to Chapters, and Paragraphs, Tables and Images are added to Sections. The document is then written to an output stream or a file. The architecture for generating the documents was divided in to three main sections. One class, the PDFGenerator, was responsible for adding elements to the document. The graphers were responsible for generating charts to be added using the Java library JFreeChart. Graphers further specialized in the sort of data they graphed for example, was the data to do with production or trade, or perhaps individual agents? Each grapher implements the class IGrapher. This required a method "createSeries" which, given an enumeration of the graph to be generate, calls the necessary method and returns you the chart.

In addition to graphing, there were also classes that dealt with general statistics, such as averages, standard deviations and so on. Finally there was the class that draws the illustrations of trade for the agents. Crossover agents are represented as orange dots. An arc between two dots illustrates a trade in that iteration. The graphics were generated using Java Graphics. Coordinates for agents were calculated and stored. If a trade occurred between two agents in that round, an arc is drawn between them. The coordinates for the start and end point were simply the coordinates of the agents, and the control point of the arc, a point through which the curve must pass, was generated using an elliptical formula.

9.4 Interface

The interface serves purely as an area in which to configure the simulation. Originally, it also displayed graphs, but as the implementation progressed, the richness of the PDF documents generated outweighed the need for the interface to deal with anything other than configuration. In order to keep the interface as simple as possible, It seemed better to remove the graphs in favour of solely offering the ability to download the simulation in PDF form.

A partial screen shot of the interface is given in Figure 9.6. Care was taken to restrict the amount of manual form entry by the user as much as possible. This was achieved through both the use of radio buttons and check boxes where applicable, and also by automatically hiding unnecessary configurations depending on the current options selected. However, for elements such as the number of agents, manual form entry was necessary, so validation was implemented in the interface to ensure that certain constraints had been met. For example,

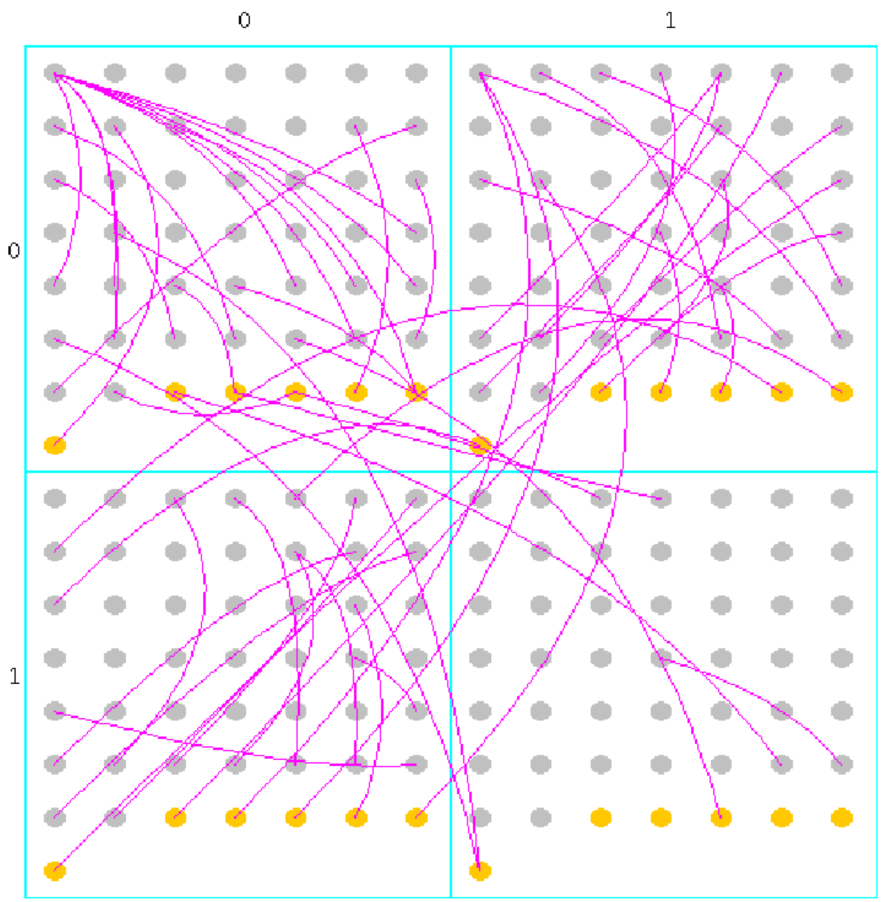

Figure 9.5: Illustration of a trade network for one iteration

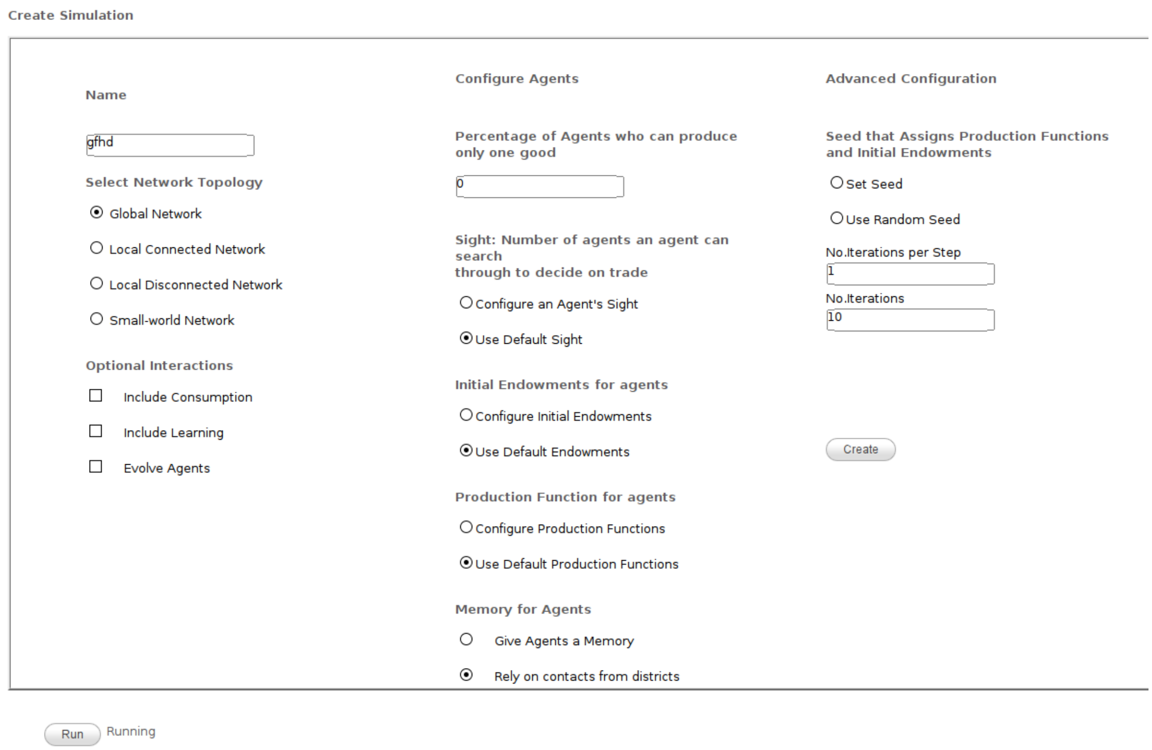

Figure 9.6: Partial Screen shot with all specific configuration hidden

ensuring that the number of agents was greater than 1, and that upper bounds were greater than lower bounds and so on. Validation by alerts is illustrated in Figure 9.8. To make the amount of configuration options less overwhelming to the user, they were hidden and shown when necessary. Comparing Figure 9.6 with Figure 9.7, it is apparent that in Figure 9.7 there are more options in the "network topology" section outlined in red.

The plain style of the interface helped create a more user friendly configuration screen. However, it is important to realise that some of the terms are ambiguous to first time users, thus a help section acting as a glossary of terms must be added.

9.5 How it works

For clarity a flow diagram has been provided (Figure 9.9) illustrating the workings of the application, which should help the reader to understand the following explanation.

The simulation is created from the web interface, with the user configuring

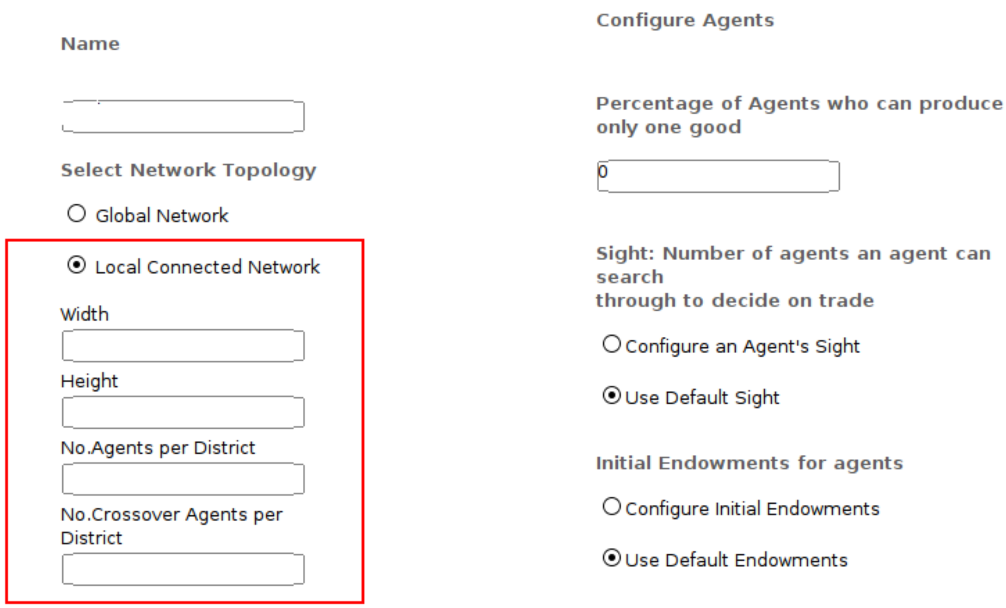

Figure 9.7: Partial Screen shot with some configuration options revealed (highlighted in red)

their simulation by turning options on and off, specifying the number of agents and districts and so on. On pressing "Create", these configuration are sent using JQuery's wrapper around AJAX - a far simpler way than the standard Javascript method, to the "ConfigServlet". This, through the interface to the simulation, SimulationServer, creates a simulation which is stored as an application level variable. The purpose of the SimulationServer is twofold. Firstly it allows an extension which distributes the simulation across various machines (that would be specified by the user) using RMI. [2] Secondly, it means that the only thing the web interface can do is create and run a simulation - everything else is closed to it.

The user is then able to select the simulation they want to run. Like creation, this request is sent to the "SimRunServlet" (again employing AJAX) where the simulation server is requested to run the simulation of the identifier (name)

[2]This was attempted at the beginning of the project, but difficulties with Tomcat's security set up caused problems with the RMI Security Manager when it came to loading Python scripts. It was abandoned due to time constraints and the subsequent hindrance that this caused with respect to the progression of the project. However, I believe it would be a computationally useful extension to add if any further work were to be done on the project. For this reason the architecture was left in place.

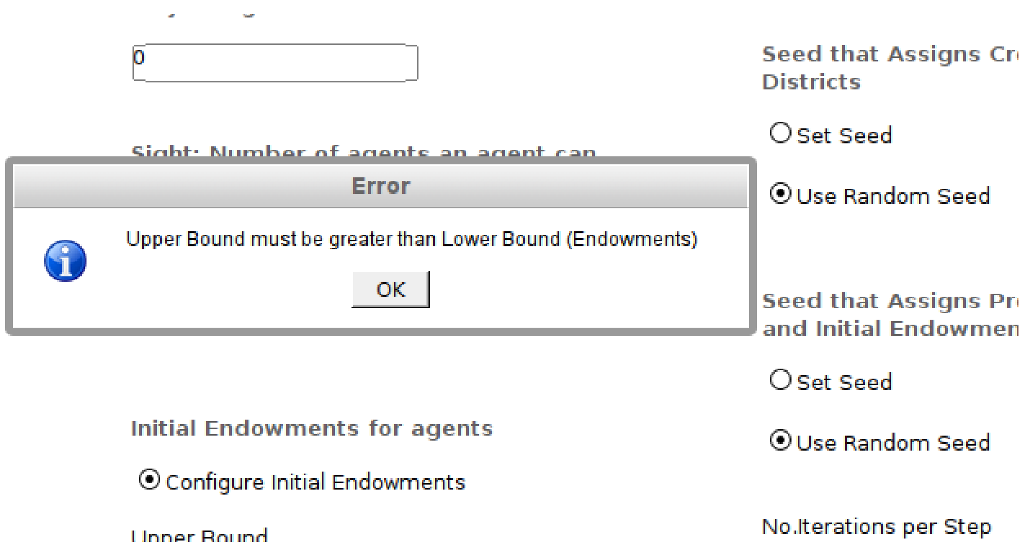

Figure 9.8: Partial Screen shot showing form validation via alerts

specified. Due to the request response architecture, and the time out on requests, the run servlet cannot be trusted to return the output of the simulation. For this reason, the servlet executes run and then the page polls the simulation every 10 seconds to see if it has completed. On receiving an all clear from the PollSimulationServlet, the page then allows the simulation to be saved (the button for saving a simulation appears). When the user clicks this button, a request is sent to the server to generate a PDF document for download.

9.6 Distributing Simulations

As mentioned, the implementation of distributing simulations using RMI was abandoned early on, due to the difficulties in getting around Tomcat's Security Manager. However, upon beginning the evaluation, it was clear that distribution was imperative in order to complete the simulations necessary. Therefore, a different way to achieve this was sought.

A web server called Jetty that provides an http server and client as well as a servlet container was utilised. This allowed Tomcat to be abandoned and replaced with Jetty. In order to distribute simulations across machines, a simple script was created that would use ssh to remote login into a machine, execute the war file and open Firefox. In this way, a list of machines was provided, and simulations could be run on each of them. I also extended the servlet that runs the simulations to run them in bulk - allowing the user to create several, and then run them sequentially.

The result was an considerably quicker collection of results and a useful

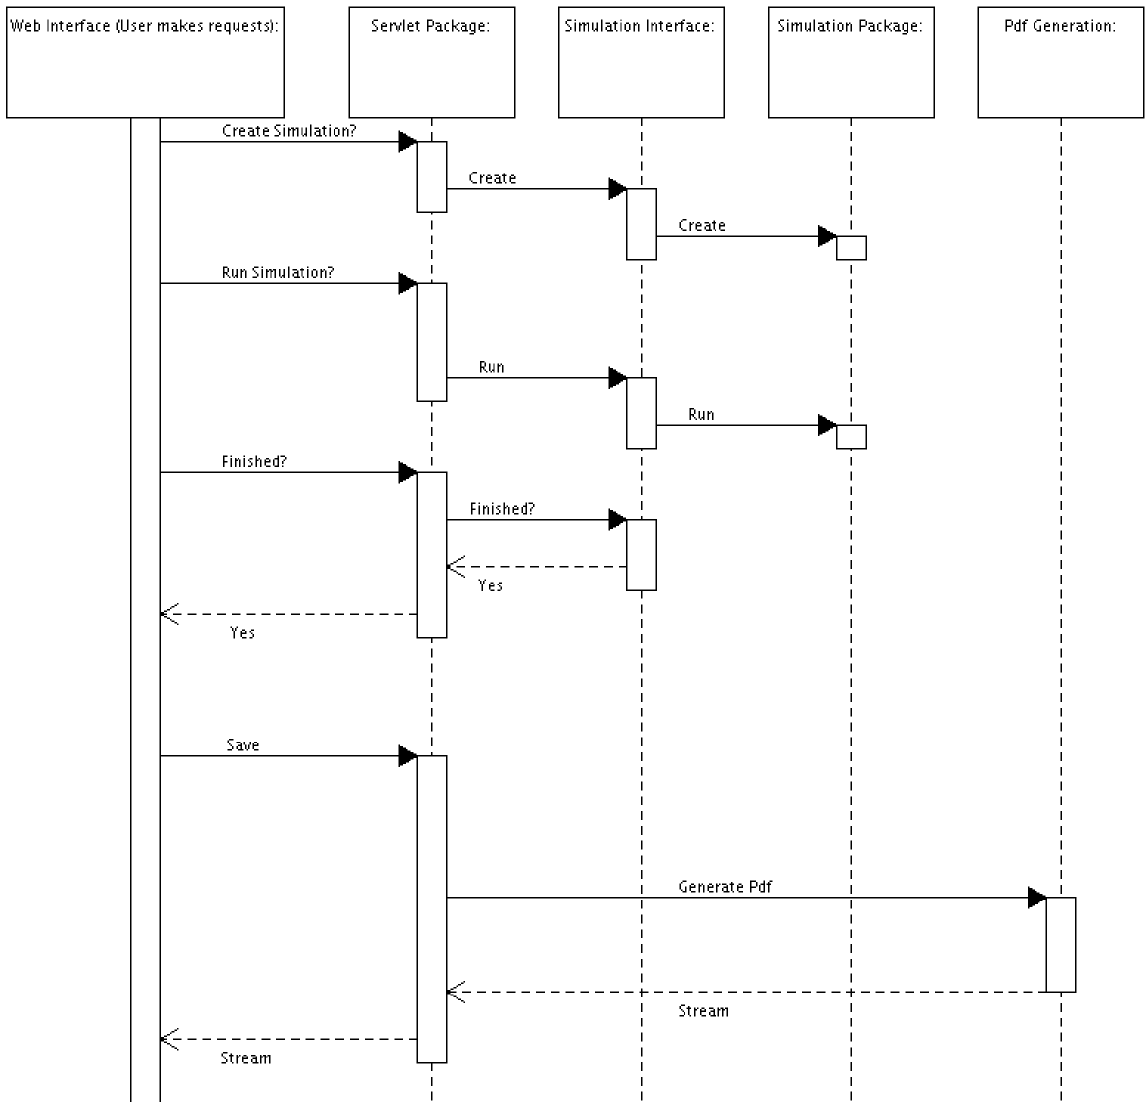

Figure 9.9: Illustration of process for using application

feature for a user who wishes to run simulations in bulk, yet quickly.

9.7 Software Development Process & Testing

The software development process followed was the iterative method. This seemed to be the most sensible option since the project was focussed on adding multiple extensions to an initial model. The core of the simulation - the initial model - was implemented first and thoroughly tested. Upon completion, the next extension was tackled, and again thoroughly tested. For every extension, the extension was designed, the model was enhanced, the interface was updated to allow for the extra configuration, and the analysis engine was updated to create graphs and tables specific to the extension.

This was not only a logical break down of the project, but also allowed for changes to extensions, for new ideas to materialise, and for the direction of the project to be steered based on how the simulations are running. In addition it meant that extensions could remain "fuzzy" until they came to being implemented, when the idea was researched, planned and evolved until it was ready to be integrated in to the simulation.

The act of testing however was extremely difficult due to the inherent non-determinism in multi-agent systems. Therefore it was imperative to exhaustively test incrementally as an extension was added in order to be certain that I was adding to a solid base. At the beginning it was possible to use JUnit in some circumstances, for instance checking that the assignment of crossover agents was being performed correctly and so on. JUnit could also be used to check the output data for undeniably incorrect evolution of the simulation, such as negative stocks of goods recorded. JUnit was possible for testing the setup and outcome of the simulation, but after this point, the non-determinism made its usefulness deteriorate. Extensive graphing of particular data (even if it was not useful in analysis of results), for instance checking that agents who were consuming always exited the economy if they didn't consume the baseline amount, was necessary as a fallback technique to counteract this flaw. In addition, when problems were found, the extensive eclipse debugger was used in order to isolate the source of the problem, and generally speaking the error could be found through manual code inspection.

Chapter 10

Conclusion and Further Work

In this chapter both successes and limitations of the model will be discussed, as well as an evaluation of implementation.

10.1 The Model

Basing the model on Wilhite's implementation had both advantages and disadvantages. The decision to base it on Pareto's idea of Pareto Superior trading was a sensible and rational application. It is realistic since in the world of business nobody ever makes decisions that could possibly make them worse off knowingly.

Specialisation was witnessed in all simulations, and as division of labour was enforced through restricting the number of agents who could produce both goods, agents experienced strategy shifts. More agents specialised in the two extreme ends of the continuum, becoming pure producers and pure traders and enhancing the credibility of the simulation. It meant that increasingly more agents could rely on the demand of the rest of the population, and, in addition, that trade was increased through the exploiting the symmetry of the Cobb-Douglas utility function, requiring agents to both increase and even out their stockpiles of goods..

The evolution of the simulation showed a strong correlation between production functions and an ability to generate wealth. This created a world in which producers were extremely wealthy. In addition, it was not the case that poor production in one of the goods was a disadvantage - on the contrary, as the percentage of agents able to produce both goods fell, this position became more advantageous. Producers could rely on demand from other agents in order to generate wealth. This meant that in one iteration, producers could engage in free exchange - exchange that they did not initiate - which led to the emergence of wealth condensation. This is an extremely realistic finding - those with money or expertise are in a position to attract more wealth. The act of traders sacrificing their goods to the producers widened the gap in wealth globally, resulting in an increase in the Gini Coefficient. The Gini Coefficient became more realistic with a decrease in the number of agents being able to produce both goods. This is an emergent characteristic since initial endowments, together with production functions, are randomly and, most importantly, uniformly distributed across the population.

The implementation of memory showed how permitting agents to remember encounters could both increase the loyalty of agents and cause the population to incur strategy shifts. Thresholds were appparent as the tensions between increased competition in the market place stemming from increasing the sight of agents and their being able to search through their memory took their toll. It led to loyalty for agents with memory enabled being overtaken by loyalty in the simulation without memory enabled. Despite the ability of agents to learn and form long-lasting trade relationships, the problem of the changing stockpiles of agents causing memory items to be outdated would be an interesting aspect to address. By implementing a form of garbage collection, where an agent can periodically remove contacts from their memory in order to reduce wasted space, this limitation could be bypassed.

Consumption further emphasised the difficult situation of traders in the new world where bankruptcy was possible. Producers were largely unaffected by this extension, but wealth distribution was increased considerably when constants were set to values which limit the amount of bankruptcy, allowing agents to survive on the edge. Bankruptcy being so common, however, is not an unrealistic thought in a model where borrowing isn't permitted. Virtually all businesses need loans to start up and in addition the large majority of families in the developed world have more liabilities than assets. With no financial security in the simulation, bankruptcy should be expected, even with low values of consumption. What was less realistic was who became bankrupt - once again it is the agents who initiated trade who drew the short straw. Their poverty in the simulation is a setback in capturing the reality of the modern world and trade. Trade does not imply poverty. In todays world, it is more the case that raw materials are produced by poor countries. It is the people who ship goods from a location where they are cheap to one where they are valued who actually make a lot of money. Trade is a massive source of income and thus the punishment of traders is something that must be addressed first and foremost.

By analysing the networks of districts when agents were to consume, dependencies and fragility in the network were witnessed, and in the simulation it was relatively simple to accurately predict which agents would exit the economy from only the network as data. This is an extremely exciting revelation. The vast majority of bankruptcy prediction is done using statistical models and neural networks. In fact, very few people have approached the problem from a perspective of interactions and network dependencies. With work and a lot of data, this could prove to be extremely useful in economics. Its usefulness stems from such an abstract approach, and although with the particular simulation implemented here its results aren't highly complex, it has a great amount of scope.

Nodes can represent anything from firms to countries, regions to industries.

To highlight the potential in this idea, it is worth giving a toy example. If a large company A goes bankrupt, this model could allow you to anticipate how companies who are dependent on A will cope. If a company has a low dependency on A, it is likely its demand will be met from somewhere else, possibly a smaller company willing to offer a good price in return for some big buys. If, however, a company is hugely dependent on A, it is likely that they are going to feel the brunt of the bankruptcy. If A happened to be part of an oligopoly (an industry ruled by a few big players), it is likely that the other big players will take on old clients of A. However, if you assessed the production capacity of these companies, and checked that the new capacity was feasible given the extra demand, and found that prices were likely to increase due to little redundancy in the network, then it is likely that deals will be made to those who offer the best price. Thus, based on the strength, capital and buying power of the dependents in the network, which can be modelled by extending the network to be a level deeper, including who the dependents supply to, it is possible to anticipate which companies will get the deals and which won't. Thus it may be possible to anticipate bankruptcies further down the chain.

However, there is another and perhaps more interesting side to this coin. If you can make a reasonable guess as to who is going to go bankrupt, then supply to their clients must be picked up by someone, just as was so for the initial bankrupted company. If you can assess the production capacity of firms, perhaps you can look at who would cope best with the new load. In turn, it could be possible to spot some good investments - companies who will be able to handle the new load and are likely to be taking this new load on would probably be good

Although work is clearly needed, computing power and access to necessary data suggest that even if this is to take time and research to make accurate, it is worth the time and research! This further highlights the importance of network structure in the study of economics, and offers suggestion that insight can be gained from looking at the economy from this more abstract perspective.

Upon implementing learning, we witnessed large increases of wealth, even above that of the Global Network (ceteris paribus) illustrating the potential benefits that can be realised through globalisation. In addition, it was suggested that perhaps it is true that becoming more globalised makes the world as a whole better off, but we saw only a small, almost negligible, change in wealth distribution. In addition, questions as to whether limitations of globalisation were implicit in its nature, or are yet to be witnessed due to the fact that we are yet to become truly global were discussed. The model not only demonstrated efficiency of free trade, but also gave insights into the justification of the gains of free trade. The study into the architecture of globalisation resembled a similar core-periphery structure to that witnessed in the worlds international trade network, and it was suggested that the extra decentralisation observed in the simulation could indeed stem from protectionist policies of governments. It was also suggested that there a strong correlation between trade deficits and the position of countries in a network, not only persistent in the results of the simulation, but also readily apparent in the real world, with countries of the G8 holding some of the largest deficits.

Through evaluating the evolutionary algorithm, it was clear that in spite of extensions creating more trade, trade was still not valued. Still a world of production rules. The following suggestion suggests some possible alterations to the model in order to remove the utter dependency of the wealth of an agent on their production functions.

10.2 Alterations to the Model

As mentioned, the poverty of traders is both unrealistic and hampering to the study of networks. If some traders could enjoy wealth akin to producers, they would make good trade partners to other poorer traders and perhaps even to pure producers which in turn would allow for longer trade chains, as traders could also offer good prices. However, the alteration to the model must be done in a way that doesn't punish pure producers for being good at what they do, but gives pure traders potential to make money in spite of their poor production function. This would require some sort of "reward scheme" since you could not just give all agents who emerge to be pure traders a random advantage.

One way of facilitating this is limiting the time agents have in which they can search the population. Perhaps one unit of time, or one iteration, could be made up of a number of units (most likely to be the sight of an agent). An agent could be allowed to trade until all of its "time units" have been used up. By employing memory of agents, agents could store encounters. At the moment, these encounters aren't ordered, but this could be altered so they are ranked by gain in utility and reliability measured in number of uses. This way, one unit of time could be used up for every potential partner an agent searches through before settling on a trade partner. This way agents could strategically search their memory, and if they are agents who have a few suitable trade partners, or high loyalty, they could make an educated guess at who is likely to offer them the best deal, and go straight to this partner. This way, agents who are successful in quickly finding trade partners are rewarded by increasing the volume of trade they can perform in a round. This acts then as the free exchange experienced by pure producers, but is not damaging to the wealth of the pure producers. In fact, it could actually increase it. It is unlikely that this would benefit the producers so much that traders are still relatively poor. There will still be some traders who are poor as they may not have reliable trade partners and as such may be in the same position as before, using up most of their units in order to find a trade partner. However, it is likely that some pure traders would gain huge benefits from this. Perhaps they can ship goods between two agents several times in one iteration. They improve their capability in moving large volumes. This could cause a positive feedback loop, whereby being able to perform multiple trades in a single iteration allows larger margins to be made, such that next iteration they are able to trade in larger quantities and thus make an even larger margin, and so on.

Alternatively, instead of time, a cost could be placed on distance. Perhaps agents can spiral outwards[6] from their position in the network, searching for a trade partner. Thus, the further out they sought, the more expensive was the transaction cost. This transaction cost can be thought of as analogous to shipping, and could be paid in Good 2. This would encourage trade to occur locally, and as such may encourage longer trade chains emerging. If the transaction cost correlated to distance however, you may prevent some long distance relationships emerging. To counteract this potential problem, transaction costs could increase with distance, but decrease with the number of times you have engaged in exchange with the same agent.

10.3 Implementation

The main issue of problems with RMI were overcome by a simple script to deploy the application on multiple machines. However, a great extension would be the implementation of a method to compare results of simulation. By outputting a CSV file of simulation data, this comparison engine could load multiple simulation files, and either perform averaging of the same sort of files or alternatively perform comparison between simulations. The most important element would be the collaboration of similar simulations into one CSV file. It would be useful to provide a method of customizing the data pulled out of the uploaded file to remove any redundancy.

Although the PDF files were useful and user-friendly, when it came to evaluation across many simulations, it proved to be both time consuming and laborious. Thus to facilitate quicker and potentially more in depth evaluation, this is a definite extension that should be made.

The trade-off between the decrease in performance of utilising Python and the speed added in the development process proved to be a good one. Extensions were quick and easy to implement, and the abstraction enforced by the divide of the scripting layer from the core facilitated better design.

The overall performance proved to be linear in terms of the number of districts, and factorial in terms of the number of agents per district. This large increase in time as the number of agents in a district increases is not a flaw in implementation, but implicit in the requirement of the model since every agent must search through every other agent in order to find a trade partner. However, these search costs can be reduced by reducing the sight of agents.

10.4 Concluding Thoughts

In conclusion, the simulation proved that with such a simple model, complex network structures can emerge that are easily likened to the real world. In addition, economic phenomena were witnessed, such as creative destruction and wealth condensation, and arguments for how they occur were presented and discussed. Agents were seen to specialise, and the extensions led to some interesting strategy shifts. Consumption, although offering insight in to the importance of network structure, redundancy and bankruptcy chains, was nonetheless a fairly unstable extension. There was a fine line between catastrophic bankruptcies and normal running of the simulation. However, a more realistic logarithmic increase in wealth was witnessed which was an improvement from the linear trend. Neoclassical economics was questioned, as prices never settled to uniform price, despite utilising the Cobb-Douglas utility function and model of Pareto superior trading, both of which have deep roots in Neoclassical economics. The production and consumption of goods caused price fluctuations, presenting a simulation more akin to an open system - with matter or products, being both created and used - as opposed to a closed static linear system.

However, the value created in the real world through moving products from one location to another was not fully achieved. Learning most increased the value as pure were enabled to become more strategic in their partner selection. Despite this shortcoming, there were many successful aspects of the simulation. In particular, the ability to produce results in the structure of the international trade network which were so close to the real world reenforced the fact that the simulation need not be overly complex in the way in which agents interact.

From conducting the simulations, although they lack realism in places such as value creation from trade, it is clear that agent-based modelling is well-suited to modelling such a complex system. Even such simplistic abstractions from reality, such as trade of knowledge, permitted astoundingly realistic results for the minimal level of detail employed.

Bibliography

[1] Allen Wilhite *Bilateral Trade and Small-World Networks* Compuational Economics 18: 49-64, 2001, Kluwer Academic Publishers

[2] Leigh Tesfatsion *Introduction* Compuational Economics 18: 1 - 8, 2001, Kluwer Academic Publishers

[3] Jie Shen *Dynamics of Human Society: Introduction to Multi-Agent System Based Research in Social Sciences* ISO report, Department of Computing, Imperial Collegen London, 2008

[4] N. Basu, R. Pryor and T. Quint *ASPEN: A Microsimulation Model of the Economy* Computational Economics 12: 223241, 1998. 223 1998 Kluwer Academic Publishers.

[5] John Cassidy, Dept. of Disputation *The Decline of Economics* The New Yorker, December 2, 1996, p. 50

[6] Camille Guèrillot *Dynamics of Human Behaviour* MSC, Department of Computing, Imperial College London, September 2005

[7] Kelvin Au *Dynamics of Human Societies: Evolution of Hierarchical Groups* Department of Computing, Imperial Collegen London, June 2005

[8] Frank Kriwaczek *Utility* Decision Analysis Lecture Notes, Department of Computing, Imperial College London, 2008

[9] Frank Kriwaczek *Decisions Involving Multiple Objectives* Decision Analysis Lecture Notes, Department of Computing, Imperial College London, 2008

[10] *Encyclopedia Britannica, 11th Edition, Jeremy Bentham* Now in the public domain, available online at *http://encyclopedia.jrank.org/BEC_BER/BENTHAM_JEREMY_1748_1832_.html* Accessed 20th January 2009

[11] *BBC Article - High oil prices hit global economies* *http://news.bbc.co.uk/1/hi/business/7421778.stm* Accessed 24th May 2009

[12] Raja Kali, Javier Reyes *The architecture of globalisation: A network approach to international economic integration* Department of Economics, Sam M Walton College of Business, University of Arkanas, May 5 2006

[13] Robin Cowan, Nicolas Jonard *Network Structure and the Diffusion of Knowledge* Journal of Economic Dynamics & Control 28, 2004, 1557 - 1575

[14] V Latora, M Marchiori *A measure of centrality based on network efficiency* New Journal of Physics, IOP Publishing Ltd and Deutsche Physikalische Gesellschaft, 2007

[15] Darrel Whitley *A Genetic Algorithm Tutorial* Computer Science Department, Colarado State University

[16] *Diagram for Gini Coefficient http://en.wikipedia.org/wiki/File:Economics_Gini_coefficient.svg* Accessed 21st May 2009

[17] *Encyclopedia Britannica, 11th Edition, Adam Smith* Now in the public domain, available online at *http://encyclopedia.jrank.org/SIV_SOU/SMITH_ADA_17231790_.html* Accessed 10th January 2009

[18] Nick Vriend *On Walrasian Models and Decentralised Economics* Research Bulletin, 1991, Vol. 3, No. 1, p. 25-37, European University Institute, Florence, Italy

[19] *The Economy as an Evolving Complex System: An emerging Paradigm* Bulletin of the Santa Fe Institute, Volume 3, No. 2, Summer-Fall 1988, p.11-13, Director of Publications: Ronda K. Butler-Villa

[20] Yaneer Bar Yam *Dynamics of Complex Systems* Addison-Wesley, 1997, Chapters 0, 1, 8, 9

[21] Allen Wilhite *Economic Activity on Fixed Networks* Handbook of Computational Economics, Leigh Tesfatsion, Chapter 20, p1014 - 1043

[22] Allen Wilhite *Self-Organizing Production and Exchange* Compuational Economics, 21: 107-123, 2003, Kluwer Academic Publishers

[23] *Understanding the Anasazi Culture Change through Agent-Based Modelling* Jeffrey S. Dean, George J. Gumerman, Joshua M. Epstein, Robert L. Axtell, Allen C. Swedlund, Miles T. Parker, Stephen McCarroll p 179 - 205 Dynamics in Human and Primate Societies By Timothy A. Kohler, George J. Gumerman

[24] Joshua M Epstein *Generative Social Science Studies in Agent-Based Computational Modeling* Chapter 1, 2006, Princeton University Press

[25] *http://support.dundas.com/OnlineDocumentation/WinChart2005/Anova.html* Accessed 31/05/2009

[26] *http://upload.wikimedia.org/wikipedia/commons/b/bd/Lorenz-curve1.png* Accessed 30/05/2009

[27] Eric D. Beinhocker *The Origin of Wealth* Chapters 1 - 9, 2007, Random House Business Books

[28] *http://www.jython.org* Accessed 29 December 2008

[29] *http://en.wikipedia.org/wiki/Agent-Based_Computational_Economics* Accessed 30 December 2008

[30] *http://jade.cselt.it/* Accessed 24 January 2009

[31] *http://www.fipa.org/repository/aclspecs.html* Accessed 24 January 2009

[32] *simon.cs.vt.edu/SoSci/converted/ANOVA/activity.html* Accessed 14 May 2009