VICTORIA UNIVERSITY OF WELLINGTON

Te Whare Wananga o te Upoko o te Ika a Maui

School of Mathematics, Statistics and Computer Science

Te Kura Tatau

PO Box 600
Wellington
New Zealand

Tel: +64 4 463 5341
Fax: +64 4 463 5045
Internet: office@mcs.vuw.ac.nz

Emergence of mutualisms in Simple Agent-Based Bartering Economies

Kieran Paton

Supervisors: Marcus Frean and Kris Bubendorfer

Submitted in partial fulfilment of the requirements for
Bachelor of Science with Honours in Computer Science.

Abstract

This report investigates emergent behaviours and trade patterns in agent-based bartering economies. A model is defined that allows production, trade and specialisation to occur within a network of agents. The model is implemented and experiments conducted to determine which factors promote mutualistic behaviour amongst all agents in a network.

Contents

Figures

Chapter 1

Introduction

This project focuses on the emergence of mutually beneficial behaviours or *mutualisms* from networks of agents engaged in the barter and production of commodities. Economic specialisation is the prime example of mutualistic behaviour considered, although other emergent patterns and interactions are also investigated.

1.1 The Problem

The paradox of mutualistic behaviour is conveyed well by the issue of specialisation.

Consider a set of agents, each of which *could* produce small amounts of everything it needs, but whose efficiency is higher if it specialises in just *one* thing. If trading is possible then it would benefit each agent (and the network as a whole) if all agents specialise and trade their surpluses equitably amongst the group. However, it does not seem to the advantage of *any one* agent to specialise in the first place – as there will be little demand for its 'product' if the other agents already have a bit of everything they need.

To gain the benefits, a simultaneous move toward specialisation seems to be required on the part of all agents involved – an improbable coincidence given the assumption of independent, self-interested agents. The problem is reminiscent of the prisoner's dilemma from game theory, in that 'something more' (besides purely rational, self-interested action) is required to induce cooperation. I shall argue in Chapter 4 however, that the situation is more accurately reflected by the game known as 'the stag-hunt' illustrated in Figure 5.1.

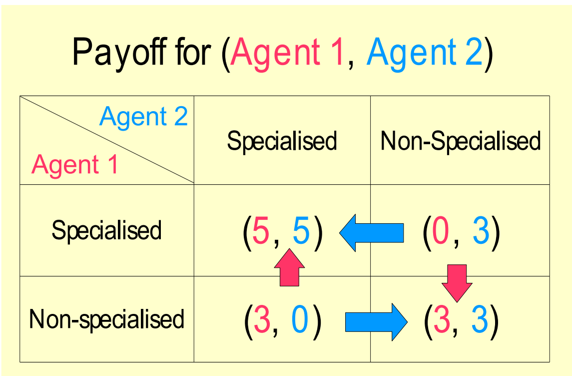

Figure 1.1: Payoff matrix indicating how the benefits of specialisation depend on the states of *both* agents involved. The arrows indicate what action a rational agent should take, *given* the state of the other agent.

The question of what features of the underlying model will induce specialisation now has motivation. And since the concept of trade is vital to that of specialisation, defining a robust model of barter is also important.

1.2 Goals

The aim of this project is to find the simplest conditions under which specialisation (as an example of a mutualism) emerges in an agent-based bartering economy.

This will involve defining and implementing a model in which agents are capable of production, barter and specialisation within a trade network. Three key questions arise:

- How will barter and production strategies affect trade patterns and how well each agent does?
- What effect does network topology have?
- What are minimal factors that will produce specialisation amongst all agents in the network?

The emergence of money from a basic bartering simulation is an issue of further interest.

1.3 Applications

While this project is based strongly in theory, its results have the potential to be applied across a wide variety of disciplines:

- The problem of resource allocation in distributed systems. One solution is the use of autonomous software agents that barter for resources – mutually beneficial behaviour would be of interest in ensuring fair alllocations.
- Explaining co-operation and altruism (mutualisms) among human groups subject to social dilemmas.
- Explaining mutualistic behaviour in biology. This could include bacterial resource usage, symbiosis or other species interactions in ecological systems.
- Investigating the origins and nature of simple economies and money – in particular how they emerged in early human societies.

1.4 Outline

The next chapter reviews the literature and previous projects on the topic. Chapter 3 outlines the basic bartering model, along with decision points and justifications for each stage. Chapter 4 extends this model to cope with production and specialisation. Chapter 5 gives a brief overview of the implementation and some basic validation. Chapter 6 gives experimental results from the basic model, demonstrating the emergence of the middleman phenomenon. Chapter 7 gives experimental results and analysis on the issue of specialisation. Chapter 8 offers some brief insights and experiments on the subject of money, and finally Chapter 9 concludes.

Chapter 2

Literature Review

This chapter provides a critical overview of the existing literature on the subject, and gives a brief description of a number of similar projects.

2.1 Literature

There is a diverse range of literature on agent-based economic modeling involving barter. Some discuss specialisation [10], network topologies [3] and the origins of money [13, 18] directly, while others give more general approaches to defining a basic bartering model [14, 15]. The brief overview given here will inform decisions and debate in later chapters.

Andrea Lavezzi's simulation of economic specialisation [10] does not explicitly model barter at all. Agents in the network decides to 'specialise' or 'despecialise' in one of two commodities, according levels of demand (based on the states of neighbouring agents). This implicitly assumes a model of barter in which commodities are traded equitably within an agent's neighbourhood. The parameters Lavezzi used include the initial states of agents, the number of neighbouring agents and threshhold levels of demand for specialisation (and despecialisation).

In general, he found that specialisation (and hence higher overall output) was facillitated by lower threshholds and larger numbers of neighbours in the network. A small concentrated set of specialised agents could also propagate a wave of specialisation throughout the network. Paradoxically, the same factors that encouraged initial specialisation caused unstable, continuously fluctuating patterns in which agents were constantly switching between specialised and non-specialised states - much like the result found in Shane Gibb's project ([7], p.33).

While interesting, Lavezzi's model has a number of limitations. First he only considers cases with two commodities, which gives a very limited set of behaviours for each agent (see section on combinatorial explosion, Chapter 4). More importantly, he considers specialisation as a black and white distinction – an agent cannot be partly specialised along some continuum. I think these limitations partially explain his results, so it will be interesting to see if the same patterns arise in my model (where those limitations will not be in place).

Phillip Bonacich's paper [3] focuses on power differences between agents in different topological networks. For instance, the agent at the hub of a star network has the power to select a trading partner and gain a better deal through competition. The outlying agents, on the other hand, have no choice but to trade with the hub. However, Bonacich is more interested in how power differences can affect the *evolution* of networks to new topologies.

In his model dissatisfied agents may modify their position in the network, changing the set of agents with whom they trade. The result was that networks with positions of unequal power tended to evolve towards stable configurations of equal power.

Exploring trade patterns in fixed networks (of different topologies) and in self-modifying, evolving networks are both interesting ideas worth pursuing further. This is supported by the economic standpoint of John Foster [5], who holds that network *structure* is crucial in determining the value of goods and services.

A variety of sources held views on the origins of money and on the relationship between money and barter.

Prendergast and Stole claim that in their model of voluntary reciprocal barter [13], trust from repeated transactions substitutes for money. In the absence of a *double coincidence of wants* (where each agent wants what the other agent has) money is usually required as a universally valued 'medium of exchange'. In their model, trust plays this role instead – agents give freely what is requested of them in the expectation that their own requests will be fulfilled at some later time. Cheating is constrained by the threat of a trading relationship being dissolved. Furthermore, they claim that in cases where reciprocal exchange is hard to enforce (due to extreme differences in valuation between agents), production of unwanted goods is higher and those goods can assume the role of money ([13], p.3).

It is generally accepted that when a double coincidence of wants does exist, bartering works well and there is no advantage to having money. Engineer and Shi however [4], argue that even in this situation money can be useful. When assymetric bargains are prevalent (due to some agents having more bargaining power), money gives fixed purchasing power across all agents and can increase the welfare of the system ([4], p.2).

Williamson and Wright ([18], pp.104-105) give a brief overview of traditional interpretations of money. These tend to focus on how the *properties* of objects (such as low storage or exchange costs, persistence and high costs of private production) can cause them to become mediums of exchange.

However, Williamson and Wright argue that private information concerning the *quality* of goods is the prime motivator for money's existence. This is corroborated by Banerjee and Maskin's argument that, even in the absence of a double coincidence of wants, barter can still take place if the *resale value* of the goods is taken into account ([2], p.956). However, if one does not know how to determine the quality of the good being sold, then one may not be able to estimate its resale value and hence run the risk of being cheated. Both these accounts conclude that money is simply a commodity with fixed quality that is universally known – and indeed Banerjee and Maskin claim to prove that goods with a very small discrepancy in quality quickly become the medium of exchange and assume the role of money in the system ([2], p.958).

Finally we come to the basic model of barter itself. Two theoretical models are presented, with some interesting insights into the nature of two agent barter.

Eric Rasmusen [14] outlines the basic bargaining problem in terms of splitting a pie of size one between two agents. His model involves each agent making alternating offers $0 \leq \theta \leq 1$ until either one agent accepts the other's offer or a fixed 'time' (number of offers) has elapsed. By applying game theoretic techniques he concludes that each agent's offer will consist of keeping the entire pie to itself - and in the final round the other agent will accept because it has nothing to lose. Rasmusen concludes that some form of discounting (decrease in value of the pie over time) or delay cost (fixed cost per offer) is necessary to ensure any result other than this simple outcome.

Ariel Rubenstein [15] presents a more mathematical formalisation of this basic model and is equally concerned with the importance of time preferences (as he calls discounting and delay costs). He also presents a model [16] in which time preferences differ between agents and each agent knows only its own – effectively leading to different bartering strategies, and more complex behaviour such as opponent identification, bluffing and reputation building.

However, this model involves splitting a pie that is neither owned nor took any effort to produce, by each agent (and hence for no offers do agents stand to make a strict lo ss). This is quite different from a situation where agents already possess commodities and then barter to maximise their net wealth (as defined by some utility function). I think this difference is crucial and could well impact on the necessity for delay costs or discounting in the bilateral bartering scenario (see Chapter 3 for more details).

This brief overview of the literature provides a useful starting point for investigations into barter, money, networks and specialisation.

2.2 Related Projects

A number of projects involving barter, trade and specialisation have been carried out or are ongoing. A brief outline of the most relevant are given here and contrasted with the approach I have taken.

2.2.1 Shane Gibbs (BIT project, 2005)

Shane Gibbs' BIT project [7] defines a model in which agents produce and trade amounts of various commodities. A utility function over these commodities defines an agent's wealth and notion of value. Agents seek to maximise their utility by simultaneously engaging in barter with their neighbours in a trade network – conducted through a process of offers and counter-offers, modified on a per round basis. Offers are transacted (the commodities exchanged) after a set number of rounds in which they are agreeable to both parties. Agents continue trading goods until a stable state is reached in which no agent can make further increases to their utility. Only complete networks, in which everyone can trade with everyone, were tried.

Specialisation was handled by picking a poorly performing agent and modifying its production by some random amount. This form of selection and improvement gradually evolves agents towards the production levels that most benefit them. This does not seem to quite fit with the 'agent-based' approach however, being more of a system imposed modification. Ideally agents should be able to modify their own production, based only on the information they have to hand – rather than any system-wide properties.

My project shares the same basic approach to the problem, but differences include a greater focus on network topologies and a more agent-based approach to specialisation.

2.2.2 The Colour Trading World

Besides Shane Gibbs' project, work by Rodolfo Garcia-Flores and Nectarios Kontoleon [6] was the closest to mine in terms of basic approach and motivation. It is primarily focused on how strategies, network topology and network positions affect an agent's performance and emergent behaviours in a network. Other similarities include restricting agents to be non-loss making, and accepting only trades they can afford. There is also a general commitment to the notion of limited private information, as opposed to agents having global access to each others states.

Offers and requests deal in colours (analogous to my commodities), which perculate through the network from 'sources' via 'merchants' to 'sinks'. This is a key difference to my model in restricting where production and consumption take place in the network.

2.2.3 MUD simulation

Grimm and Mitlohner's MUD (multi-user domain) simulation [8] is based around the network game-playing genre. Both human *and* artificial agents may engage in economic interactions with each other in a simulated environment, and a form of money is used rather than strict bartering *per se*. However, some features such as the way 'deals' are conducted may still be pertinent. A deal consists of partners meeting bilaterally to negotiate terms, through an unlimited series of offers and counter-offers until agreement is reached. If one partner cannot fulfil the terms then their rating from the viewpoint of the other agent is decreased.

Although quite different from my model, the notion of bilateral barter is intriguing (as opposed to simultaneously engaging with multiple agents), along with the notions of trust implied by memory of past interactions with other agents.

2.2.4 Trade Network Game

Leigh Tesfatsion's Trade Network Game forms part of his webpage and project on "Agent-Based Computational Economics: Growing Economies from the Bottom Up" [17]. It is a simulation of a virtual economic world, capable of exhibiting individual behaviours, trade patterns and social welfare in decentralised market economies. It allows for agents to cooperate and defect, evolve over time, choose trading partners and learn from experience. It can simulate a variety of markets and deal in commodities ranging from labour time, earnings and money, to apples and oranges – in other words it can handle barter [11].

However, many aspects of the simulation involve more complex economic behaviour than desired for this project.

2.2.5 OPTIMAES

OPTIMAES is the 'Open Project To Investigate Money And Economic Systems' by Phil Jones and Hilan Bensusan [9]. It consists of a Python simulation for modeling a variety of economic features including barter, production, consumption and currency. Agents are situated in a social network that may be fully or locally connected and they intend to expand the simulation to cope with modeling small-world and scale free networks. Statistics collected by the simulation relate to both individual agents and the network as a whole.

While having many similarites to my approach, the simulation is geared more toward observing strictly economic effects rather than emergent behaviours or mutualisms *per se*.

2.3 Summary

The literature and projects described tend to focus on a single issue (such as barter, network topology, specialisation or investigation of money) and hence tend to design their models and simulations narrowly around it. Little work has been done on the interplay between these factors, and how a variety of complex patterns and behaviours can emerge from a simple simulation (apart from [7, 6]). The main contribution of this project will be to see how such behaviours and patterns (like specialisation) can emerge from a basic bartering simulation.

Chapter 3

Basic Model of Barter

This chapter gives a summary of the key points and decisions taken when defining the basic model of barter. The model consists of a collection of agents able to possess and trade in a number of commodities. These concepts will be explained and defined in depth, with justification given for the design choices made. Simplicity of both model and implementation were key considerations in the process.

Many of these decision points offer interesting asides, experiments and points of investigation that will be noted and could be examined in a larger work.

3.1 Agents

A **trading agent** is an independent, self-interested entity consisting of a *state* and a *behaviour*. For the sake of brevity, trading agents will henceforth be referred to as agents.

3.1.1 State

An agent's state consists of a finite **inventory** of commodities, detailing how much of each commodity is possessed by the agent.

A **commodity** is a tradeable item possessed in continuous, non-negative amounts by an agent. Continuous amounts were chosen, because the bartering process is much simpler if commodities are infinitely divisible – as agents may then set and modify their offers precisely, according to a given valuation. The expected behaviour would be the same as for discrete commodities in large numbers. It would still be an interesting aside to investigate bartering involving small numbers of discrete commodities and hence test the following hypothesis:

Hypothesis 3.1 *More plentiful (or infinitely divisible) commodities will act as a currency for scarcely occurring, indivisible commodities.*

Bartering involving debt (negative values for commodities) would be a further avenue to explore, but adds complexity to the model by removing some constraints on an agent's actions.

3.1.2 Utility

The **utility** of an agent defines its net wealth and notion of value. It is defined as a function over the commodities possesed by an agent, and so in this case is a direct consequence of the agent's state. The following function is one of many possibilities:

$$Utility = \sum_{i=1}^{n} \log(c_i + 1) \tag{3.1}$$

where $\{c_1, c_2, ..c_n\}$ represents the inventory of the agent. One is added within the logarithm to ensure a *positive* value whenever the commodity amounts are positive.

This function was chosen because it reflects a simplistic view of supply and demand in a market. As figure 3.1 shows, an agent's utility (for a fixed sum of commodities) is maximised by having equal portions of each commodity. It is also maximised by having as much of each commodity as possible. Directly related to utility is an agent's **valuation** of a particular commodity. This can be interpreted as the extent to which an increase in that commodity will increase the agent's utility. As a consequence, agents tend to value scarcely occurring commodities highest (the ones they tend to possess least of) and commonly occurring commodities lowest (the ones they tend to possess most of) – in other words, supply and demand effectively determine wealth and value in the market.

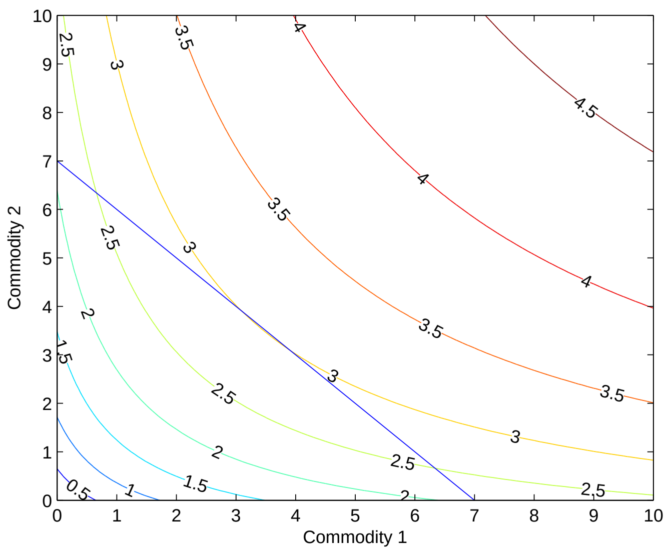

Figure 3.1: Contour plot of utility function for a 2-commodity agent. The blue diagonal line indicates the utility given a fixed sum of commodities (7 in this case).

This is certainly not the only utility function that makes sense. Other possibilities include taking the intrinsic properties of different commodities into account, as opposed to just their relative amounts (giving a notion of intrinsic value). More interestingly, the notion of resale value could be added to more accurately assess the future or potential worth of an agent's inventory. The following hypothesis takes inspiration from Banerjee and Maskin [2]:

Hypothesis 3.2 *The addition of resale value to an agent's utility function will facillitate the emergence of money-like commodities – those commodities with a high ratio of resale value to intrinsic value, will be the ones more likely to resemble currency.*

Also, it might be interesting to model scenarios where agents have different systems of valuation (different utility functions) and even different commodity sets. However, in the interests of ensuring maximum *potential* cooperation among agents, I decided to make these features universal. The rationale behind leaving these possibilities for now is to keep the model as *simple* as possible until the basic experiments and phenomena have been observed.

As an example of the topics discussed so far, the state and utility of 3 agents in a 2-commodity simulation might be displayed as follows:

	Commodity 1	Commodity 2	Utility
Agent 1	0	10	2.398
Agent 2	10	0	2.398
Agent 3	5	5	3.584

3.1.3 Behaviour

An agent's behaviour consists of making **offers** of commodities to another agent or agents. To simplify the model it is assumed (and enforced) that all agents are honest and rational:

- By **honest**, I mean that an agent will only propose offers that it is capable of fulfilling.
- By **rational**, I mean that an agent will only propose offers that will increase its utility if transacted.

The presence of trading agents that are *not* honest in this sense would be an interesting addition to the model, allowing for notions of trust, bluff, deceit and so on to play a role (see Prendergast and Stole [13]).

Hypothesis 3.3 *Agents not focused on immediate profit (not 'rational' in this sense) may reveal situations in which a short-term loss is required to obtain an eventual long term gain.*

In both cases though, the agent's range of possible actions would be expanded and significant complexity added to the model.

As for the actual make up of offers, two main models were considered:

- First there is the classic notion of 'splitting a pie' between two agents as illustrated in figure 3.2. The two agents decide which commodities they wish to trade in and then barter until they agree on a *ratio* to split the goods. Presumably the actual amount of goods to be exchanged must also be settled.
- The second idea is illustrated in figure 3.3 and involves each agent presenting an offer (consisting of a **commodity** and an **amount**) to its trading partner. The combination of the two offers must be agreed upon by both parties for the specified amounts to be exchanged. This model was ultimately chosen, but the reasoning was not straightforward.

The first model had the appeal of needing only a single number (the ratio) to represent offers between two agents. Besides having the potential to simplify the bartering process, some nice analyses also came out of it (see "The problem of two agent barter"). This initial simplicity was misleading however, given that agents still needed some mechanism for signaling acceptance of a trade and furthermore had to agree which commodities to trade in and how much to transact after the ratio was agreed.

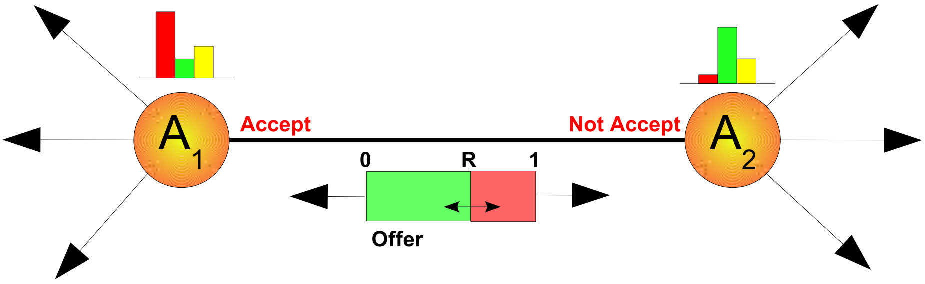

Figure 3.2: Diagram showing first way of modeling offers – 'splitting a pie' to determine a ratio.

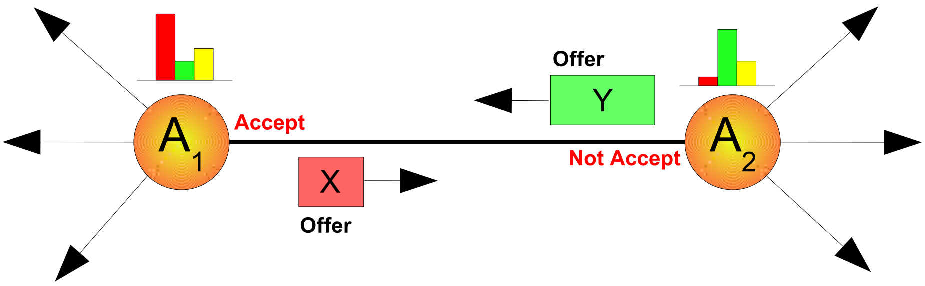

Figure 3.3: Diagram showing second (and eventually adopted) way of modeling offers – specifying independent offers and amounts.

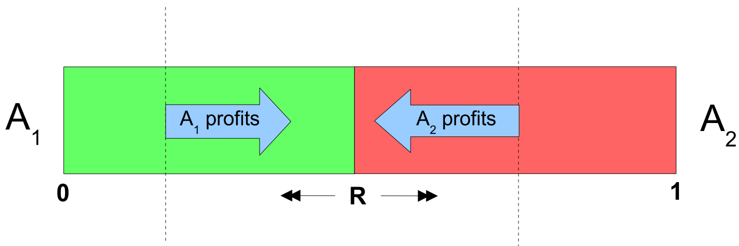

Figure 3.4: Diagram showing minimum profit levels of agents involved in bilateral barter. The current ratio R indicates a trade of mutual profit.

The problem of two agent barter:
Recall Rasmusen and Rubenstein's arguments that some form of discounting or delay costs are required to make sense of the two agent bargaining problem. Without these they argue, an agent would claim 'all of the pie' and on the last round the other agent would accept because they have *nothing to lose* [14]. Figure 3.4 illustrates a key difference between this view and the model defined here, in that agents already in possesion of the commodities they are bargaining over certainly *do* have something to lose.
There is a level below which agents will not make a profit and hence not wish to continue trading past (see diagram). I argue that the risk of termination substitutes for the discounting factors proposed by Rasmusen and ensures that fair, rational trading may take place. Note that while the pie-splitting model revealed this analysis, the conclusion generalises to two agent bartering in either model.

The last point proved to be the flaw in the approach, as changing the amounts once the ratio has been fixed can impact significantly on whether that ratio is still profitable to the agent.

Proof: Say Agents 1 and 2 have the following inventories:

	Commodity 1	Commodity 2	Utility
Agent 1	10	0	2.398
Agent 2	5	10	4.190

and that they have agreed to trade Commodity 1 for Commodity 2 at a ratio of 0.2 (that is, 0.2 of the total amount traded goes to Agent 1).
If the total amount traded is 1 unit, then the resulting commodities and utilities are:

	Commodity 1	Commodity 2	Utility
Agent 1	$10 - 0.8 \times 1 = 9.2$	$0 + 0.2 \times 1 = 0.2$	2.505
Agent 2	$5 + 0.8 \times 1 = 5.8$	$10 - 0.2 \times 1 = 9.8$	4.296

An increase of about 0.1 in utility for both agents. However, if the total amount traded is 10 units (at the same ratio) then the resulting commodities and utilities are:

	Commodity 1	Commodity 2	Utility
Agent 1	$10 - 0.8 \times 10 = 2$	$0 + 0.2 \times 10 = 2$	2.197
Agent 2	$5 + 0.8 \times 10 = 13$	$10 - 0.2 \times 10 = 8$	4.836

A *decrease* in utility for Agent 1 and an increase for Agent 2 – this shows that two agents cannot simply agree on a ratio and sort out the amount to be traded later, as the two concepts are inextricably linked. Unless one fixes the amount to be traded at some arbitrary value (say the unit amount) then this method would seem to be of little further use. The problem is graphically illustrated in figure 3.5 – note how the ratios at which trade becomes profitible for each agent shift significantly when the total amount traded is changed.

There are several advantages to the second method however, including greater applicability to the multi-agent scenario. Because offers are simply sent to other agents (rather than being directly linked with them by a ratio), it is easy to imagine agents sending and receiving multiple *interacting* offers and hence modelling direct competition. Furthermore, the message-like nature of offers means that a special form of 'offer' may be used to signal the acceptance or rejection of the current trade.

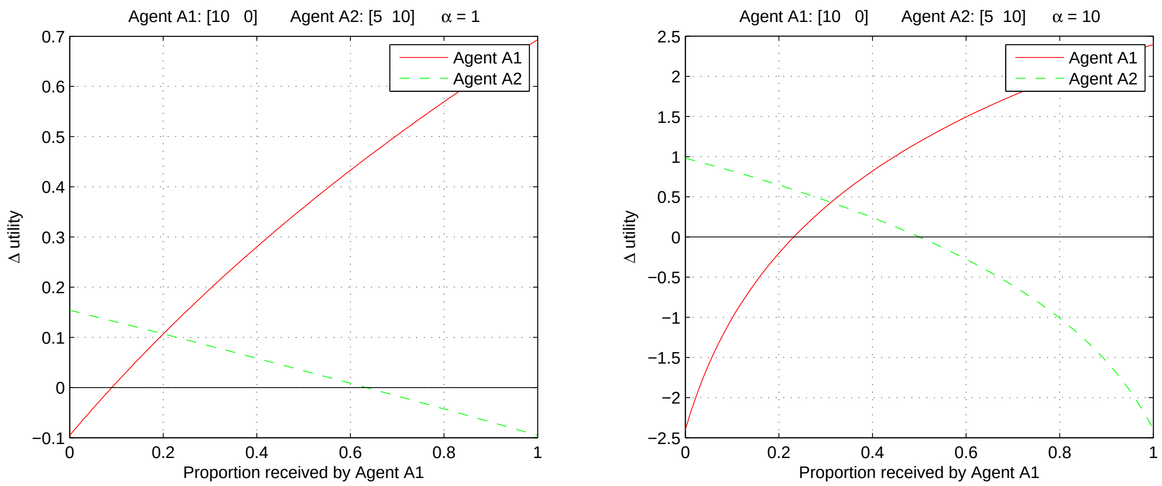

Figure 3.5: Graphic illustration of problems with 'pie splitting' method. The y-axis shows the change in utility for each agent that would result from the trade, according to the ratio given on the x-axis. α gives the total amount traded (1 on left, 10 on right).

3.2 Trade networks

A **tradelink** is a connection between two agents that allows them to interact – that is, trade by exchanging offers. It consists of a current offer (possibly empty) and a boolean variable for each of the two participating agents – the boolean denoting whether or not they accept the current offers (as shown previously in figure 3.3).

A **trade network** can be represented as a graph where the nodes are agents and the edges are trade links. One way to implement some notion of trust (or at least past satisfaction), would be for agents to weight these edges on the basis of successful transactions and use this information in determining who it trades with in the future (and how).

A **trade** is a sequence of correlated offers between two agents that may or may not result in agreement. It is started by each agent making an initial offer, after which the agents make counter offers in response to both the current trade, and others it is participating in. The counter offers must be in the same commodity as the initial offer (to ensure some consistency) and the process is ended by either both agents being in an accept state (signaled by an accept message) or by one agent sending a reject message. The former results in a transaction (see below) and the latter resets the offers and the accepts to null, for instance allowing an agent to start trading in another commodity. Hence trades can be seen as threads or 'conversations' between agents that may be terminated by either party – a crucial requirement as outlined earlier in the "the problem of two agent barter".

Competition is modeled by an agent *simultaneously* engaging in multiple concurrent trades with other agents (on other tradelinks) and hence their decisions may be influenced by offers from these other parallel 'threads'.

Hypothesis 3.4 *Competition between agents should lead to significant differences in trading behaviour between the two agent versus the three or more agent bartering scenarios.*

A **transaction** is the actual exchange of commodities resulting from a successful trade and takes place *as soon as* both agents on a tradelink are in an accept state. After a transaction, the offers and accepts on a tradelink are set to null to allow a new trade (in potentially different commodities) to take place.

An alternative possibility would be the notion of 'accept by degrees', whereby an agent signals its 'happiness' about the current offers and instantly receives a corresponding proportion of the amounts specified in the offer. Besides doing away with the traditional 'back and forth' notion of barter, I suspect this method would suffer the same problems of ratios versus amounts encountered in the pie-splitting scenario (see proof, section 3.1.3).

3.3 Barter Strategies

Hence we can see that an agent's bartering strategy should consist of deciding for each tradelink:

1. What commodity to trade in.
2. An amount of that commodity to include in the initial offer to another agent.
3. When to accept the current offers.
4. When to reject the current offers.
5. If not accepting or rejecting current offers, how to modify its own current offer.

To maintain the notion of an agent-based model, the information on which an agent makes these decisions should consist only of its own state and current offers – along with any calculations that can be made from these. The most useful information includes:

$$\begin{aligned}
\{c_1, \ldots, c_n\} &= \text{agent's inventory} \\
u &= \text{agent's utility} \\
\Delta u_{current} &= \text{change in the agent's utility that would result from} \\
&\quad\ \text{the offers on the current tradelink} \\
a_{current} &= \text{whether the offers on the current tradelink have been} \\
&\quad\ \text{accepted by the other agent} \\
\Delta u_{pending_1}, \ldots, \Delta u_{pending_p} &= \text{changes in the agent's utility that would result from} \\
&\quad\ \text{pending offers on each of the other tradelinks}
\end{aligned}$$

Here **pending** offers are those already accepted by another agent. Even this restricted set of information would be a lot for a simple learning algorithm (for instance reinforcement learning) to operate on – especially given the range of possible actions at each stage. An evolutionary approach would be another alternative, but would probably require further parameterisation of the problem. Investigating these options would be a challenging project in itself – different barter strategies suggest whole new avenues to explore, including how strategies interact and perform relative to each other and the strategic environment (with possibilities of seeing the nice, nasty, robust and evolutionary stable strategies of Axelrod [1]). Evolving and learning these strategies (or even just specifying them), would be non-trivial however.

For the experiments in this project, I decided to define a default bartering strategy as follows. An agent should:

1. Choose the commodity it has most of initially, but after a random period without transactions, change to the next highest commodity it posseses and so on.
2. Initially offer one twentieth of its stockpile of the selected commodity.

3. If the current offers would result in a profit (increase in utility) and this increase is greater than or equal to that of all other trades currently pending, then accept.
4. If the opposing offer is less than a certain fixed amount, then reject (stops long pointless trades over infinitely divisible commodities).
5. If not accepting or rejecting current offers, then either harden (decrease) or soften (increase) own current offer by some amount. If the other agent is in an accept state, then decrease offer by ten percent. If the other agent is not in an accept state, then increase offer by ten percent.

Although some of the values above seem quite arbitrary, this is a reasonable strategy for a self-interested agent to follow. The justification for 1 is that agents would naturally prefer to get rid of commodities they have most of (under the given utility function) – but should be prepared to switch to a different commodity if there is *no demand for it*. Step 3 states that if the current offer is the best available, then accept it. Only when it is not the best available (or it is not profitable), should the offer be modified by hardening or softening it depending on whether the other agent has accepted or not.

The merit in this approach is that fixing the barter strategy allows the effects of other factors to be assessed.

3.4 Structure of Trading Rounds and Cycles

A **trading round** consists of each agent having the opportunity to make an offer (either initial, counter, accept or reject) to every other agent that it is connected to by a tradelink. Hence every tradelink is utilised twice in a trading round, once in each direction.

To ensure complete fairness, the order of selection for tradelinks and direction (that determines when it is an agent's 'turn') is chosen by random permutation. This means that over the course of at least two trading rounds, the same agent may get to make consecutive offers on a tradelink before the other agent has a chance to respond. In the interests of fairness therefore, the accept states of an agent are automatically set to false as soon as any change is made to an opposing offer.

Trading rounds continue until a fixed number have passed without any transactions occuring (and this number should be suitably large to allow for possible changes of preferred trading commodity by agents). A sequence of trading rounds terminating in this manner is called a **trading cycle**.

Various other structures for trades and trading rounds were considered. An agent could simply pick another agent by some method and then barter to a standstill, before selecting another agent the next round and so on. The advantage of this approach is that an agent could more easily disciminate who they traded with – useful if implementing notions of trust and past trading success.

However, I decided to go for the concurrent, multi-agent approach outlined above due to its greater ability to model direct competition in the bartering process.

3.5 Summary

A basic model of barter has been outlined that allows agents to trade commodities with each other in a network, and measure their success by a utility function. Many alterantive paths and possibilities were also described that could form a basis for a larger investigation of the topic. The next chapter outlines how the model can be expanded to cope with production and specialisation.

Chapter 4

Modeling Production and Specialisation

This chapter expands on the basic barter model given in the previous one. New mechanisms and structure are defined, allowing agents to control the production of commodities and to change this production over time.

4.1 Modeling Production

To model specialisation, some mechanism is needed to determine how much of each commodity an agent starts with. This should reflect both the risks and potential benefits that specialisation brings.

4.1.1 Capacity and Production Functions

A key aspect of production is that agents should be restricted in the total amount they can produce – modeling the simple reality of limited available resources.

This can be achieved by assigning each agent a fixed **production capacity**, which may be imagined as representing limited time or materials. In line with previous decisions to ensure consistency in the simulations, all agents will have the same production capacity. Agents with different production capacities could be used to model more complex networks involving distinct suppliers, distributors and consumers, much like the sources and sinks mentioned in [6].

This fixed capacity however, must be fully **allocated** amongst the different commodities by a partition. For instance, an agent with production capacity 3 may allocate it as [1 0 2] amongst three commodities. This implies that an increase in the production of any one commodity will result in an equivalent decrease in production of some other commodity (or commodities).

If the entire production capacity is allocated to a single commodity (as in [3 0 0]), this represents a fully **specialised** state. If production capacity is allocated evenly amongst all commodities (as in [1 1 1]), this represents a fully **non-specialised** state. Any other allocation represents a degree of specialisation between these two extremes.

The allocated capacity is then converted to actual commodity amounts by a **production function**, intended to reflect the benefits (or drawbacks) of specialiasation. As such, a general exponential function should be capable of exhibiting the required behaviour:

$$commodity_i(p_i) = k(b^{p_i} - 1) \tag{4.1}$$

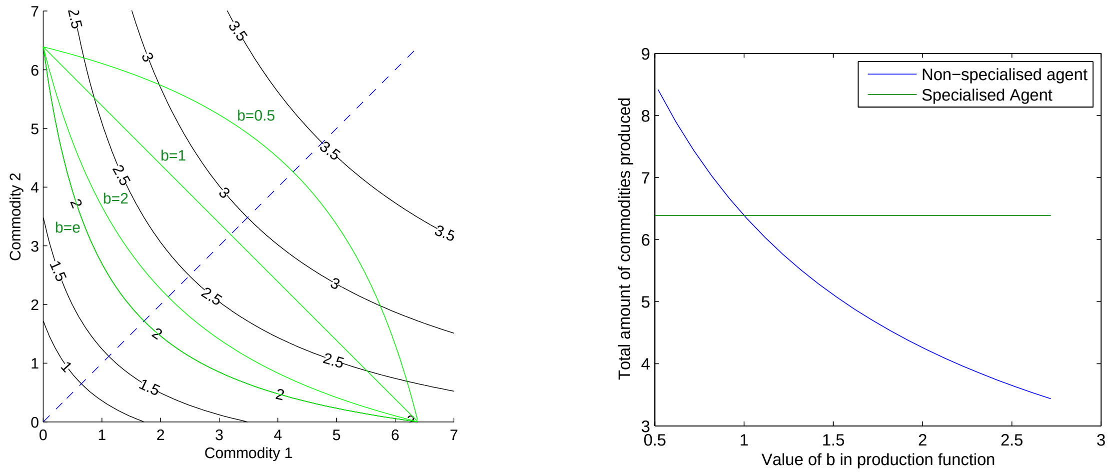

Figure 4.1: Production functions for different values of b, overlaid on a contour plot of utility (left). The green lines show the commodities produced for all allocations of production capacity 2 (A1 is specialised, A2 non-specialised). On the right is a plot showing the benefits of specialisation for different values of b.

where $[p_1 \dots p_n]$ is the allocation of production capacity $p_{total} = \sum_{i=1}^{n} p_i$. One is subtracted to ensure that zero allocated capacity results in zero commodity produced, and k is simply a constant to ensure comparability between production functions.[1] Different values of b in the above formula reflect different levels to which specialisation is a benefit or a drawback. This is illustrated in the case of two commodities by the diagrams in figure 4.1. To summarise:

- when $b > 1$ specialisation is beneficial, as a greater sum of commodities is produced when in a specialised state (devoting effort to one thing results in increased efficiency).
- when $b = 1$ specialisation is neutral as the sum of commodities produced is the same regardless (it reduces to a linear function).
- when $b < 1$ specialisation is detrimental, as a greater sum of commodities is produced when in a non-specialised state (two commodities may be easier to produce together than separately).

This suggests that different production functions might be applied to different commodities (or relationships between commodities) to reflect their inherent production costs. While intriguing, it is the mutual *benefits* of specialisation being investigated here and hence a production function with $b > 1$ was chosen for all agents and commodities – the specific value of b to be justified in the next section.

4.1.2 Risk

Initially I had the production function above with $b = e$, which so happened to be the inverse of the utility function. This meant that while the sum of commodities produced was greater in a specialised state, the *utility* of the agent was the same across all levels of specialisation (without trading) – note how the production function with $b = e$ shown in figure 4.1 perfectly overlays the contour where utility equals two.

[1] $k = \frac{e^a - 1}{b^a - 1}$, where $a =$ the total production capacity

This was a problem, because it did not reflect the inherent risk involved in specialisation. Intuitively, specialisation should be beneficial only if one can find a trading partner who wants your goods. Otherwise, it should be worse than remaining in a non-specialised state – as you end up with lots of stuff you don't want and wanting lots of stuff you don't have. The following payoff matrix, illustrating the stag-hunt from game theory helps to clarify this point:

	Take Risk (Specialise)	No Risk (Generalise)
Find Trading Partner	5	3
No Trading Partner	0	3

There are two main ways of incorporating this notion of risk into the model:

- encode risk individually, as some form of penalty based on an agent's past rate of failure and success. Effectively agents learn which states carry risk and which do not.

- encode risk into the model by *assuming* that specialised states are high risk and non-specialised states are low risk.

While the first method may be the more rigorous, I suspect it would only confirm the assumption given in the second – which is far easier to implement. Changing the production function to have $b = 2$ means that a non-specialised agent will have a higher utility before trading than a specialised agent (see figure 4.1). The latter however, will have a greater sum of commodities and hence higher utility after trading – if it can find someone to trade those commodities with.

This gives an adequte encoding of the notion of risk expressed above and ensures that specialisation is not trivially beneficial.

4.2 Modifying Production

The ability to change allocations of production capacity is crucial in allowing agents to specialise and despecialise over time. There are question of what structure these changes should be made under and what startegies should be used by agents to effect this change.

4.2.1 Structure

Two main structures present themselves as possible contexts for this change:

- A continuous dynamic model of interspersed barter and production. For instance, agents might be able to convert gains from trading into additional production capacity, which could immediately be allocated to produce more commodities, in the middle of the barter process.

- An iterative process of **production rounds**. Agents produce according to their current allocation of production capacity, carry out a complete trading cycle, modify their allocations accordingly and then repeat the process. Agents have the same production capacity each round.

While the first method sounds interesting in terms of its dynamic and reactive potential, it would be difficult to implement and more importantly, difficult to analyse. The second method has the advantage of having a clear point at which to collect and display statistics (at the end of each production round). Iterations also define a notion of time, which could be useful for implementing features like consumption or commodity decay (potentially useful for investigating money).

For these reasons I chose the iterative approach, with agents having their commodities reset to zero each round (representing total consumption). A useful analogy may be a farmer who grows some combination of crops in a year, takes them to a market place to trade and revaluates what to grow the following year based on his or her prior trading success.

4.2.2 Production Strategies

I decided that agent's should modify their production on an individual basis, using only information directly accessible to them (their recent performance, result of their trading in the previous production round and so on). This reinforces the agent-based approach taken in the rest of the project and contrasts strongly with the evolutionary-style selection and improvement of badly performing agents applied in [7].

To keep strategies simple and to reflect the notion that agents are greedy (focused on immediate profit), I decided on the following general pattern for production strategies:

1. Each production round, increase the production (allocated capacity) of a selected commodity by some small amount and decrease production of a random other commodity by the same amount.

2. Every r rounds, if the resulting utility after trades has increased over this period then keep the curretnly selected commodity – otherwise select a new commodity to increase.

Effectively, if an agent is doing well it keeps doing what its doing – otherwise it does something else. The point of step 1 is to reduce the range of actions of an agent to 'specialising in commodity i', greatly simplifying the agent's decision. A random commodity was chosen to decrease production of, because this was simpler (and equivalent over time) to redistributing the increased capacity over all other commodities. In step 2 the selected capacity is reevaluated r times to allow time for the effects of the last change to become apparent. A small random element may also be added to r, to ensure that all agents do not change simultaneously and so have some chance to react to each other's decisions.

Note that a whole range of strategies may be implemented by specifying how (or on what basis) the new commodity is selected in step 2. These include:

- random selection of a commodity – the simplest strategy with the least assumptions.

- selection on amounts left after trading – if an agent has very little of a commodity after trading, increase its production as there must be high demand for that commodity.

- selection on *change* in amounts from trading – an agent should increase production of the commodities it is selling most of, as there must be some demand for these commodities.

Hypothesis 4.1 *The random selection strategy outlined above should be sufficient to promote specialisation, but if not, one of the above* **heuristic** *strategies will be required to induce it.*

4.3 Summary

In this chapter, mechanisms and structure allowing production and specialisation to occur amongst agents in a trade network have been defined and justified.

Chapter 5

Implementation

This chapter outlines the implementation process for my simulation. Existing packages and software are considered, the structure of my implentation given and the veracity of the simulation demonstrated by two basic experiments.

5.1 Consideration of existing packages and software

While I considered software packages and open source software as starting points for my simulation, I discarded them for a number of reasons.

Ascape is a typical example of an agent-based modeling package, reviewed in a paper by Miles Parker [12]. While it is reasonably easy to set up and run a model in Ascape, understanding the details of how it works and exactly what code is being executed is more difficult. For a project like mine where simplicity and transparency of assumptions is critical, this could prove to be a major disadvantage. It could also be questioned whether the time taken to learn the details of the package is worthwhile for the one off implementation of a model.

OPTIMAES and the Trade Network Game (described in section 2.2) were open source projects (with readily available code in Python and C++ respectively), presenting the opportunity to take and modify these simulations for my own use. However, some fairly significant changes would be needed to reduce them to a level of simplicity suitable for my simulation and even then might not fucntion exactly as desired.

It was quicker and easier to design and implement my simulation from scratch in the Java programming language. This was possible because the limited scope of my investigations did not require huge or particularly complex simulations. It also allowed for full knowledge of the internal workings of the program and relative ease of implementation due to my familiarity with Java. Had large or long running simulations been a greater focus, then more efficient languages would have been considered.

5.2 Description of my Implementation

My program consists of three main classes:

Agent: This class contains the state (commodity amounts), utility and production functions, production allocations, as well as functions specifying barter and production strategies. It includes all the behavioural functionality and encapsulated (private) information, allowing an agent to interact in a network context.

Tradelink: Handles all trading interactions between a pair of agents, including transmitting offers between them, handling the acceptance mechanism and exchanging commodity amounts in a transaction.

Simulation: Consists of a set of Agents (nodes) and Tradelinks (edges) that together define a particular network topology. Handles all the structuure and turn based aspects of trading rounds, cycles and production rounds.

A small class is used to represent offers, and a Stats class used to gather the results on each experiment.

5.3 Validating the Implementation

To provide confidence in the validity of the simulation, some basic experiments were conducted to see if results matched expectations. The statistics were averaged over 100 repititions of the experiment – see section 6.1 for a detailed explanation of what they mean and how they are collected.

5.3.1 Two agent experiment

This experiment shows the basic barter model defined in Chapter 3, working for the case of two agents dealing in two commodities.

Initial Conditions

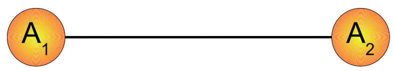

Figure 5.1: Trade network with 2 agents.

Initial states of agents are:

	Commodity 1	Commodity 2	Utility
Agent 1	0.000	10.000	2.398
Agent 2	10.000	0.000	2.398

Agents 1 and 2 are connected by a tradelink, as shown in figure 5.1.

Hypothesis 5.3.1 *Since there is an exact coincidence of wants, the agents will trade at a rate of equal exchange until they have equal amounts of each commodity.*

Results

The resulting commodities and utilities of agents after trading has finished are:

	Commodity 1	Commodity 2	Utility	Net Increase	% Increase
Agent 1	5.123	4.877	3.583	1.185 (0.000)	49.4%
Agent 2	4.877	5.123	3.583	1.185 (0.000)	49.4%

Social Welfare rose from 4.796 to 7.166 (49.4% increase)

Details of tradelink:

	Number of Transactions	Amounts exchanged	
A1 ↔ A2	14.0	A1 → C1:5.123	C0:5.123 ← A2

Analysis

Trade continued in this case until both agents possessed approximately equal portions of each commodity. The reason that exactly equal amounts were not reached is tied to the fixed size of agent's initial offers, which at some point will no longer yield a profit. Once this near equilibrium point has been achieved, the interests of the two agents become almost exactly the same and they have no further incentive to trade with each other.

The figure in brackets for the net increase gives an indication of the variance due to trading order (see next section), and a value of zero indicates that trading order was immaterial in this case.

As expected, the exchange rate of commodities was exactly equal.

5.3.2 Three agent experiment with production

This demonstrates that the simulation can cope with more than two commodities or agents and shows how commodities are produced from a given allocation of production capacity (according to a production function with $b = 2$).

Initial Conditions

Allocations of production capacity 3:

	Commodity 1	Commodity 2	Commodity 3	Specialisation Level
Agent 1	3.00	0.00	0.00	3.00
Agent 2	0.00	3.00	0.00	3.00
Agent 3	0.00	0.00	3.00	3.00

The resulting commodities produced are:

	Commodity 1	Commodity 2	Commodity 3	Utility
Agent 1	19.086	0.000	0.000	3.000
Agent 2	0.000	19.086	0.000	3.000
Agent 3	0.000	0.000	19.086	3.000

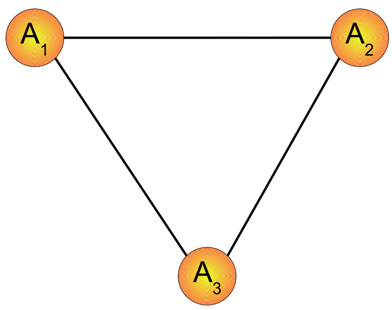

Figure 5.2: Trade network with 2 agents.

Agents are fully connected by tradelinks, as shown in figure 5.2.

Hypothesis 5.3.2 *In this symetric situation, the agents will again end up with approximately equal amounts of each commodity at equal rates of exchange.*

Results

The resulting commodities and utilities of agents after trading has finished are:

	Commodity 1	Commodity 2	Commodity 3	Utility	Net Increase	% Increase
Agent 1	6.389	6.330	6.354	5.987	2.986 (0.035)	99.5%
Agent 2	6.349	6.402	6.322	5.986	2.986 (0.037)	99.5%
Agent 3	6.349	6.354	6.410	5.992	2.992 (0.034)	99.7%

Social Welfare rose from 9.000 to 17.965 (99.6% increase)

Details of the tradelinks:

	Number of Transactions	Amounts exchanged	
A1 ↔ A2	10.4	A1 → C0:6.348	C1:6.330 ← A2
A1 ↔ A3	10.3	A1 → C0:6.352	C2:6.344 ← A3
A2 ↔ A3	10.3	A2 → C1:6.354	C2:6.332 ← A3

Analysis

Agents traded equitably as expected, until they possesed roughly balanced portions of each commodity.

5.4 Summary

This chapter has given a brief overview of the implementation process, including justification for the approach taken, a descripition of the structure of the simulation and some basic experiments to give confidence in its veracity.

Chapter 6

Experimental Results – the Middleman

This chapter presents results from the basic bartering model defined in Chapter 3. The most important result is how agents exploit positions of power in particular topological networks – the so called 'middleman phenomenon'. An overview of the statistics involved is followed by a number of experiments and a general analysis explaining the significance of the results.

6.1 Statistics

After a trading cycle has been completed, various statistics are collected and used to analyse the results of the simulation.

6.1.1 Diagnostic information

Each agent's resulting state (commodities and utility) is displayed and compared to their initial state. A summary of their performance is given by their percentage increase in utility over the course of the trading cycle. It is still important to look at the net difference in utility, as agents that start with a very low utility will naturally find it easier to increase it (due to utility being a logarithmic function).

The combined utility of all agents (the *social welfare* of the system) and its change over the course of a trading cycle is used to assess the performance of the network as a whole.

The details of each tradelink such as the number of transactions carried out on it and the relative amounts of the commodities exchanged, gives a useful indication of how much the interests of the two agents involved are compatible or in conflict.

6.1.2 Repetition of Experiments

Note that due to the random order of selection for tradelinks, experiments will not necessarily produce the same results on repetition. Hence for all the statistics mentioned above, the mean of 100 repeated experiments (with the same initial conditions) was taken.

As the order in which agents submit their offers for each trading round is the only variable on repetition, the variance (or standard deviation given here) has the potential to show how order affects an agent's power and performance in a network.

Standard deviation is given for the net increase in utility only (since all other statistics are in some way related), and is shown in parentheses.

6.2 Experiments

A number of experiments follow, each consisting of a statement of initial conditions, a hypothesis, the results of the experiment and a brief analysis. Note that the production and specialisation aspects of the model are not under consideration here, so the initial commodity amounts are simply specified for the purposes of each experiment.

6.2.1 3 Agents, 2 Commodities (Competition)

The purpose of this experiment is to investigate how two agents with competing interests define trading behaviour in a fully connected network.

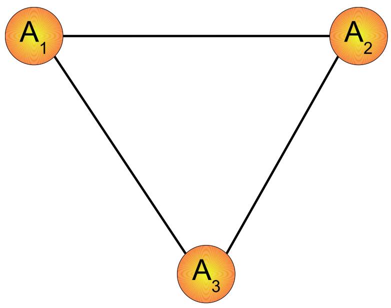

Figure 6.1: Trade Network for 3 Agents.

Initial Conditions

	Commodity 1	Commodity 2	Utility
Agent 1	10.000	0.000	2.398
Agent 2	10.000	0.000	2.398
Agent 3	0.000	10.000	2.398

Agents 1, 2 and 3 are fully connected by tradelinks as shown in Figure 6.1.

Hypothesis 6.2.1 *Since Agent's 1 and 2 have similar interests, little trade should pass between them and they should do worse than Agent 3 (who may trade with either).*

Results

	Commodity 1	Commodity 2	Utility	Net Increase	% Increase
Agent 1	6.856	2.695	3.365	0.967 (0.026)	40.3%
Agent 2	6.868	2.686	3.364	0.966 (0.024)	40.3%
Agent 3	6.276	4.620	3.704	1.306 (0.012)	54.5%

Social Welfare rose from 7.194 to 10.433 (45.0% increase)
Details of tradelinks:

	Number of Transactions	Amounts exchanged	
A1 ↔ A2	0.1	A1 → C0:0.016	C1:0.008 ← A2
		A1 → C1:0.007	C0:0.013 ← A2
A2 ↔ A3	7.4	A2 → C0:3.135	C1:2.687 ← A3
A1 ↔ A3	7.4	A1 → C0:3.141	C1:2.693 ← A3

Analysis

As expected, very little trade on average passed between Agents 1 and 2 – effectively defining a restricted trade network topology with Agent 3 as the middleman, even though Agents 1 and 2 had the *opportunity* to trade with each other. In having the power to choose which of these two agents to trade with, Agent 3 exploits the competition in the network and gains a correspondingly higher increase in utility.

Note that this can also be viewed in the context of supply and demand, as Agent 3 also possesed the scarcest commodity and was the sole supplier of it to the market. Supporting these views, Agent 3 extracted a better rate of exchange than the other two agents.

There was minimal variation on repetition, indicating that order was not a great factor in this experiment.

6.2.2 3 Agents, 2 Commodities (Middleman 1)

The purpose of this experiment is somewhat different to the last. In this experiment, the two agents separated by the middleman ideally *want* to trade with each other, but are restricted by the topology of the network. The middleman Agent 2 is given only a token amount of one commodity to encourage it to start trading.

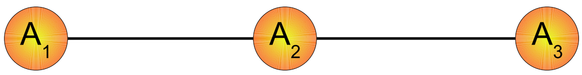

Figure 6.2: Trade Network for 3 Agents (with middleman).

Initial Conditions

	Commodity 1	Commodity 2	Utility
Agent 1	10.000	0.000	2.398
Agent 2	1.000	0.000	0.693
Agent 3	0.000	10.000	2.398

Agents 1 and 2 are connected by a tradelink, as are Agents 2 and 3 - but NOT Agents 1 and 3 (as shown in Figure 6.2).

Hypothesis 6.1 *Trade will flow through the network between Agents 1 and 3, until these agents have nearly balanced amounts of each commodity. Agent 2 will exploit its position in the middle to gain a significant profit from this trade.*

Results

	Commodity 1	Commodity 2	Utility	Net Increase	% Increase
Agent 1	6.554	1.165	2.696	0.298 (0.168)	12.4%
Agent 2	3.070	3.069	2.743	2.050 (0.522)	295.8%
Agent 3	1.376	5.766	2.689	0.291 (0.146)	12.1%

Social Welfare rose from 5.489 to 8.128 (48.1% increase)

Details of tradelinks:

	Number of Transactions	Amounts exchanged	
A1 $\leftrightarrow$ A2	8.9	A1 $\rightarrow$ C0:3.446	C1:1.165 $\leftarrow$ A2
A2 $\leftrightarrow$ A3	12.6	A2 $\rightarrow$ C0:1.376	C1:4.234 $\leftarrow$ A3

Analysis

In contrast to the previous experiment, Agents 1 and 3 are not competing for access to the middleman's commodity – instead they are trying to pass their commodities through this agent to the other side. As the exchange rates show, Agent 2 is able to demand a significant premium for this service and gain the corresponding increases in utility.

It is a valid question why trading stops when Agents 1 and 3 still possessed a significant imbalance of commodities. One possibility is that Agent 2 has gained a balance in commodities, and stops trading because it is 'satisfied'. The diagram in figure 6.3 tells a different story though. The initial states of agents are marked by circles on a commodity graph, with changes in commodities over the course of a trading cycle represented by a line (the slope of which indicates the exchange rate). Note that Agents 1 and 3 could not continue trading much further at the current exchange rates before they start going downhill in utility space. Hence it is them, not the middleman who determines when trading stops.

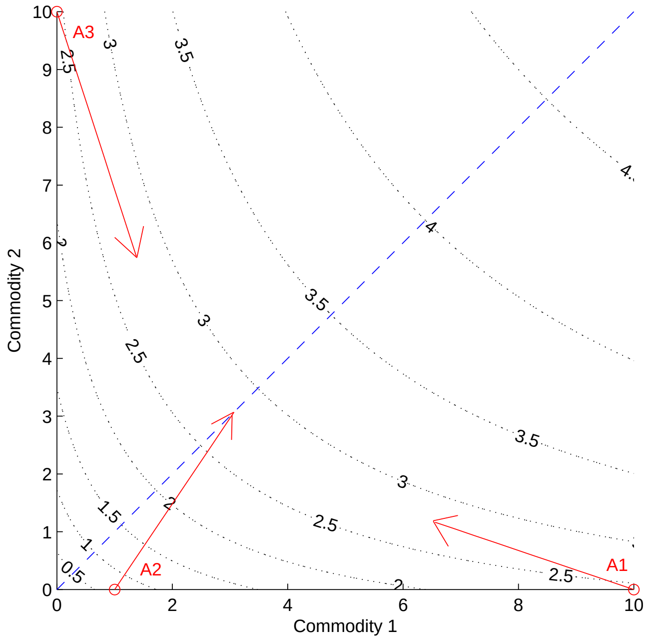

Figure 6.3: Agents' change in commodities from trade, overlaid on contour plot of utility.

The standard deviation is significantly higher in this case for all agents, but particularly for the middleman. This suggests that trading order could have a significant effect on an agent's natural position of power in a network.

6.2.3 3 Agents, 2 Commodities (Middleman 2)

To reinforce the results of the previous analysis, this experiment is identical except that the middleman starts with a balanced amount of commodities.

Initial Conditions

	Commodity 1	Commodity 2	Utility
Agent 1	10.000	0.000	2.398
Agent 2	1.000	1.000	1.386
Agent 3	0.000	10.000	2.398

Agents 1 and 2 are connected by a tradelink, as are Agents 2 and 3 - but NOT Agents 1 and 3 (as shown in Figure 6.2).

Hypothesis 6.2 *In light of the previous analysis, the middleman should still be able to make significant gains in trade until halted by the reluctance of the other two agents to trade further.*

Results

The resulting commodities and utilities of agents after trading has finished are:

	Commodity 1	Commodity 2	Utility	Net Increase	% Increase
Agent 1	7.614	0.906	2.640	0.242 (0.225)	10.1%
Agent 2	2.503	2.722	2.436	1.049 (0.730)	75.7%
Agent 3	0.882	7.372	2.617	0.219 (0.212)	9.1%

Social Welfare rose from 6.182 to 7.693 (24.4% increase)

Details of tradelinks:

	Number of Transactions	Amounts exchanged	
A1 $\leftrightarrow$ A2	6.2	A1 $\rightarrow$ C0:2.386	C1:0.906 $\leftarrow$ A2
A2 $\leftrightarrow$ A3	7.0	A2 $\rightarrow$ C0:0.882	C1:2.628 $\leftarrow$ A3

Analysis

Although the middleman (Agent 2) starts with balanced commodities, it still wants to maximise them at the expense of the other two agents. The large difference in commodity amounts between Agent 2 and the other agents makes this possible and profitable for all parties – though far more so for the middleman.

The middleman is unable to gain the same level of commodities (or hence utility) as in the previous experiment – and in fact ends up with a lower utility than the other two agents. This might partially be explained by the high level of variance in the results, showing that trading order has a particularly significance here. It also seems to imply that starting with a balanced level of commodities restricts the middleman's desire to trade, and reduces the flow of commodities in the network.

6.3 Significance of Results

These results explore ideas of power in a network and how agents exploit them for profit. The emergence of the middleman behaviour is important, in that it demonstrates that the agents are acting in a self interested, profit seeking manner.

Two distinct situations were found, suggesting that power in a network does not solely stem from an agent's position in that network's topology (the main focus of Bonacich's work [3]). The notion of competition is equally important, where two agents with exactly the same interests (high levels of the same commodity) must compete for trading access to a third party. While it may seem that this third party is acting like a middleman (in that no trade passes between the competing parties), commodities are not actually flowing through it – rather it is acting like the sole distributor of a product to two consumers. The 'restricted topology' is a side effect arising from the simulation as opposed to a given condition, as in the true middleman scenarios.

Neither case exhibits particularly mutualistic behaviour. While all agents do profit, the middleman profits at the expense of the other two agents.

More work could be done on the importance of order in aiding or establishing positions of power in the network. Investigating fixed orders in which one agent always makes the first offer in a trading round would reveal if any advantage or disadvantage is incurred. The variance in results on the middleman experiments suggests that these positions of power are highly subject to the order of trading, which would be a good future point of investigation.

6.4 Summary

The exploitation of positions of power by agents is an important emergent behaviour, resulting from the basic barter model applied across a variety of initial conditions and topological networks. It demonstrates almost the exact opposite of the mutualistic behaviour under investigation here, and highlights the problem of getting self-interested agents to cooperate for mutual benefit.

Chapter 7

Experimental Results – Specialisation

This chapter presents the experimental results on specialisation. It shows that attaining specialisation was reasonably straightforward in the two agent, two commodity case, but that difficulties arose when simulations with more agents and commodities were considered. The need for a heuristic strategy is argued for, rather than just the simple random strategy with the least assumptions.

7.1 Statistics

Since the basic barter model was not the main focus of these experiments, a different set of statistics was used to analyse their results. Change in production is the most important feature, so the resulting production allocations are given and compared to the initial allocations for each agent.

Production rounds were continued until either a stable pattern had emerged or it was fairly clear that no such stable pattern would occur.

The change in production over the course of the experiment is illustrated by a graph, showing the level of specialisation for each agent at each production round. Specialisation level is simply taken to be the maximum amount of capacity allocated to any one commodity. Production capacity is always set to the number of commodities in the experiment, so the minimum level of specialisation is always one and the maximum level is equal to the number of commodities.

For the 3 agent experiments, it was also useful to plot the actual production allocations for each commodity and each agent, to gain a clearer understanding of the problem.

7.2 Experiments with 2 agents and 2 commodities

These experiments were carried out to see if specialisation developed in the simplest network, consisting of two commodities and two agents connected by a tradelink. A production function with $b = 2$ is assumed (unless otherwise stated) and the simple random selection production strategy (outlined in section 4.2.2) is used by both agents.

7.2.1 Basic Experiment 1

The purpose of this experiment is to see if two agents can specialise and remain specialised, starting from *non*-specialised states.

Initial allocations

	Commodity 1	Commodity 2	Specialisation Level
Agent 1	1.00	1.00	1.00
Agent 2	1.00	1.00	1.00

Hypothesis 7.2.1 *The two agents should specialise in two different commodities and remain in that specialised state.*

Results

Allocations after 500 production rounds:

	Commodity 1	Commodity 2	Specialisation Level
Agent 1	1.99	0.01	1.99
Agent 2	0.01	1.99	1.99

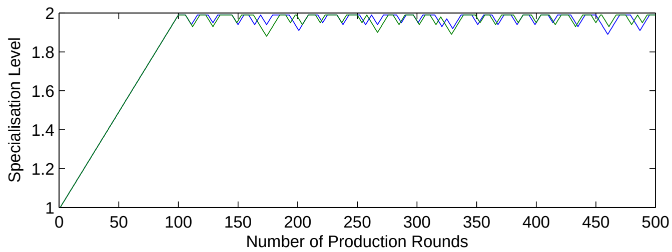

Figure 7.1: Specialisation levels of two agents over 500 production rounds.

Analysis

The two agents simultaneously and immediately begin to increase production of different commodities. Once specialisation had been attained, small fluctuations are seen – this is because the utilities of the agents ceased improving at this stage, causing them to experimentally select another commodity to increase production of. This would result in a drop in their utility however, prompting them to change back and hence keeping the two agents in a stable state of near specialisation.

7.2.2 Basic Experiment 2

The purpose of this experiment is to make sure that the two agents did not specialise in two different commodities by luck. This time each agent starts specialised in the *same* commodity, which should be highly deleterious to both agents.

Initial allocations

	Commodity 1	Commodity 2	Specialisation Level
Agent 1	2.00	0.00	2.00
Agent 2	2.00	0.00	2.00

Hypothesis 7.2.2 *At least one of the agents will return to a non-specialised state and start specialising in the other commodity.*

Results

Allocations after 500 production rounds:

	Commodity 1	Commodity 2	Specialisation Level
Agent 1	2.00	0.00	2.00
Agent 2	0.02	1.98	1.98

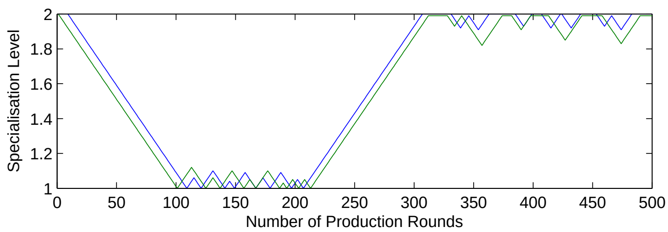

Figure 7.2: Specialisation levels of two agents over 500 production rounds.

Analysis

One agent did begin to despecialise, and interestingly the other agent soon followed. This should be no real surprise considering the disadvantage of being the only specialised agent in a network (as pointed out in section 4.1.2). A period of fluctuations followed, but ultimately the two agents again began to specialise in two different commodities as before.

7.2.3 Basic Experiment 3

The purpose of this experiment is to confirm that the production strategy in use is not biased towards specialisation itself, but is simply acting in the best interests of the agent.

A production function with $b = 0.5$ was used in this experiment, which makes non-specialisation more beneficial than specialisation (as explained in section 4.1.1).

Initial allocations

	Commodity 1	Commodity 2	Specialisation Level
Agent 1	2.00	0.00	1.00
Agent 2	0.00	2.00	1.00

Hypothesis 7.2.3 *The agents will quickly fall into a non-specialised state and remain there.*

Results

Allocations after 500 production rounds:

	Commodity 1	Commodity 2	Specialisation Level
Agent 1	0.98	1.02	1.02
Agent 2	0.96	1.04	1.04

Analysis

As expected, specialisation levels fell steadily for both agents, ending up in a stable state of non-specialisation.

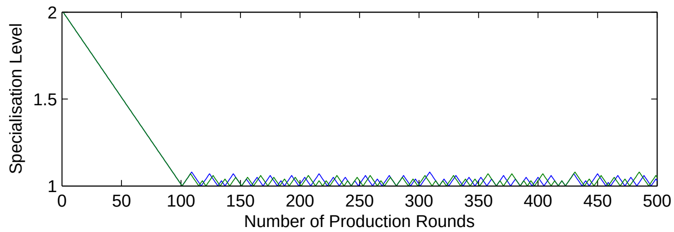

Figure 7.3: Specialisation levels of two agents over 500 production rounds (for production function b=0.5).

7.2.4 Discussion

These results are important in demonstrating that two self-interested agents can move towards a state of mutual specialisation and remain there, using a very simple strategy based only on their recent performance (utility changes). The major fluctuations between specialised and non-specialised states reported by Gibbs and Lavezzi [7, 10] were not seen.

One major reason for the difference with Lavezzi's results was in having a *continuum* of specialisation. This allowed any fluctuations to be localised around one end of the continuum, rather than between each extreme (as would be the case in a two state system with only fully specialised and fully non-specialised states).

Figure 7.4 shows a game theory style payoff 'matrix' for each agent, but for a continuum of specialisation states (compare with the stag-hunt matrix in figure 5.1). Hence payoffs (the average utility after trading) are represented by a contour graph. The lowest point for an agent is where that agent is specialised and the other is not (blue) – the highest point is where both agents are specialised (red). The blue circles plot the change in specialisation levels for each agent over the course of an experiment. Note how the agents move steadily uphill in 'payoff' space and then randomly walk around the area of the maxima trying to find a better solution.

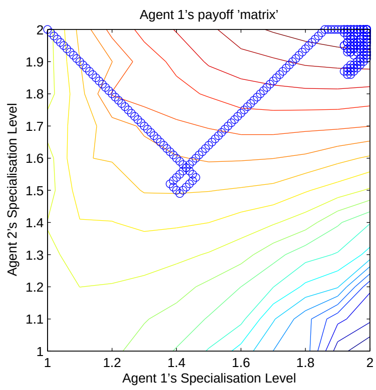

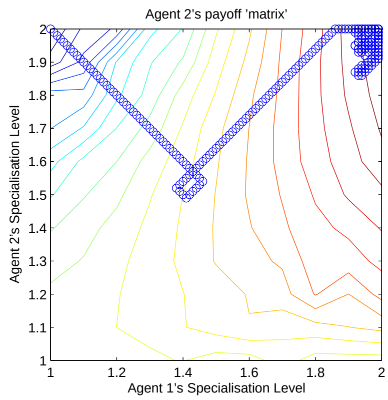

Figure 7.4: Diagram showing change in specialisation level for two agents over the course of several hundred production rounds, plotted on a contour graph of the agent's 'payoff'

7.3 Experiments with 3 agents and 3 commodities

In this section I tried (unsuccessfully) to generalise the results just gained to networks with 3 commodities and 3 agents (fully connected by tradelinks). Again the basic random selection strategy described earlier and a production function with $b = 2$, are used.

The actual allocations of each agent to each commodity are graphed over the course of the experiments, as this gives a better understanding of what is going on in the three agent scenario.

7.3.1 Experiment 1

The basic experiment is repeated for the case of three agents and commodities. Agents start non-specialised.

Initial allocations

	Commodity 1	Commodity 2	Commodity 3	Specialisation Level
Agent 1	1.00	1.00	1.00	1.00
Agent 2	1.00	1.00	1.00	1.00
Agent 3	1.00	1.00	1.00	1.00

Hypothesis 7.3.1 *Agents should specialise and remain that way, as for the two agent case.*

Results

Allocations after 4000 production rounds:

	Commodity 1	Commodity 2	Commodity 3	Specialisation Level
Agent 1	0.16	2.31	0.53	2.310
Agent 2	1.60	0.47	0.93	1.600
Agent 3	0.31	0.40	2.29	2.290

Change in production allocations over the course of the experiment shown in figure 7.5.

Analysis

Somewhat surprisingly, no clear or stable pattern of specialisation emerged in this case. Although Agents 2 and 3 attained reasonable levels of specialisation around production round 2000 (as figure 7.5 shows), these were not maintained or held by all agents at any one time. A possible explanation is that the additional agent made the chance of all the agents' actions coinciding for mutual benefit far less likely.

The amount of time needed to see patterns emerge is also significantly greater than in the two agent case.

7.3.2 Experiment 2

This experiment is designed to test the intuition expressed in the previous analysis, by seeing the effect of holding one agent fixed in a fully specialised state.

Initial allocations

	Commodity 1	Commodity 2	Commodity 3	Specialisation Level
Agent 1 (fixed)	3.00	0.00	0.00	3.00
Agent 2	1.00	1.00	1.00	1.00
Agent 3	1.00	1.00	1.00	1.00

Hypothesis 7.3.2 *Fixing one agent will decrease the coincedence in action required and promote a more stable state of specialisation in the other two agents.*

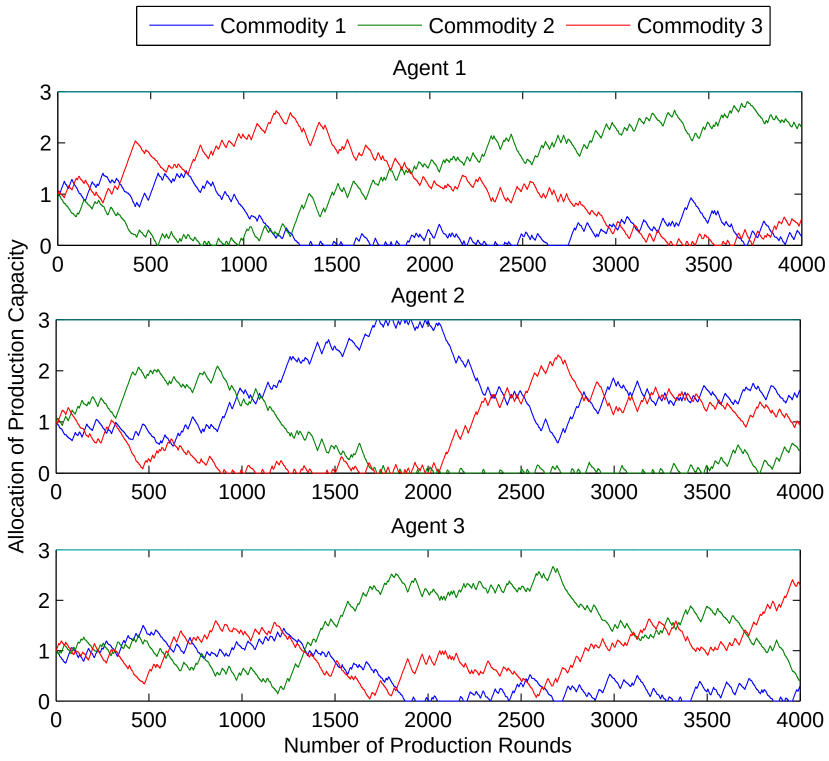

Figure 7.5: Allocations of three agents over 4000 production rounds (Experiment 1).

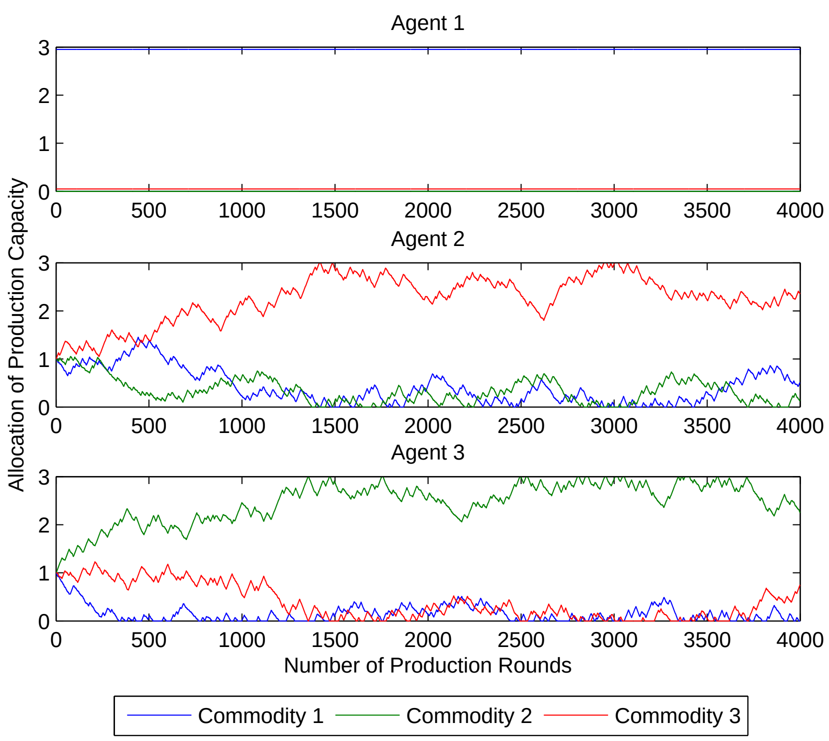

Figure 7.6: Allocations of three agents over 4000 production rounds (Experiment 2 – one agent fixed).

Results

Allocations after 4000 production rounds:

	Commodity 1	Commodity 2	Commodity 3	Specialisation Level
Agent 1 (fixed)	3.00	0.00	0.00	3.00
Agent 2	0.49	0.16	2.35	2.35
Agent 3	0.00	2.24	0.76	2.24

Analysis

Figure 7.6 shows a significantly greater move towards stable specialisation than in the previous experiment. This suggests that the number of agents acting and reacting to each other has an important effect on the emergence of mutual specialisation. Although fluctuations are greater than in the two agent model, this may partially be explained by an agent's random selection strategy applied to three commodities. If an agent in a specialised state experiments by selecting a different commodity (that then decreases its utility), it may then take several goes to get back to the 'right' commodity (that increases its utility again).

7.3.3 Experiment 3

I also decided to test the effect of having two specialised agents initially (and one non-specialised) to see if the agent's specialisation 'environment' could lead to the emergence of stable patterns.

Initial allocations

	Commodity 1	Commodity 2	Commodity 3	Specialisation Level
Agent 1	3.00	0.00	0.00	3.00
Agent 2	0.00	3.00	0.00	3.00
Agent 3	1.00	1.00	1.00	1.00

Hypothesis 7.3.3 *The two specialised agents will promote specialisation in the third agent.*

Results

Allocations after 4000 production rounds:

	Commodity 1	Commodity 2	Commodity 3	Specialisation Level
Agent 1	2.56	0.08	0.36	2.560
Agent 2	0.26	2.68	0.06	2.680
Agent 3	0.19	0.52	2.29	2.290

Analysis

The presence of already specialised agents did seem to promote the emergence of stable specialisation in all three agents. While Agent 2 also reacted to the short supply of commodity 3 by increasing production of it (production rounds 500 - 1500, see figure 7.7), it eventually relented in favour of Agent 3.

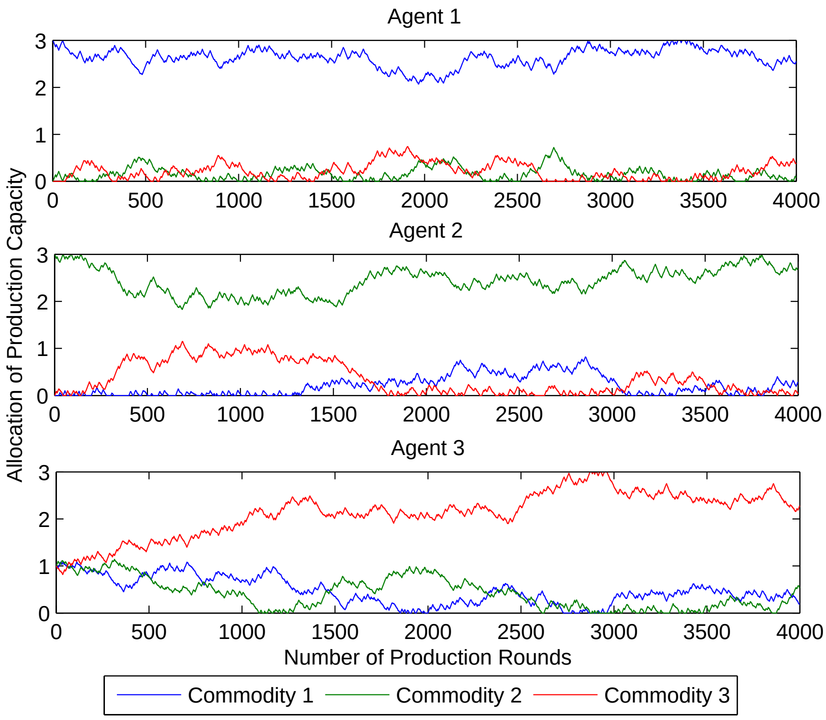

Figure 7.7: Allocations of three agents over 4000 production rounds (Experiment 3 – all agents start specialised but 1).

7.4 The problem

The experiments in the previous section reveal a problem for mutual specialisation (and mutualisms in general) in networks involving more than two agents. The range (and combinations) of actions available to agents expands significantly in these cases, suggesting that a heuristic to promote specialisation is required.

7.4.1 Combinatorial Explosion

In the two agent case there are only four possible combinations of actions that the agents may be taking, two of which are mutually beneficial:

Agent 1 action: Increasing production of	Agent 2 action: Increasing production of	Mutually beneficial?
Commodity 1	Commodity 1	no
Commodity 2	Commodity 1	yes
Commodity 1	Commodity 2	yes
Commodity 2	Commodity 2	no

In the analogous case for 3 agents and 3 commodities, there are 27 possible combinations only three of which are mutually beneficial to all agents. The odds of happening on this combination by chance have thus reduced significantly (from $\frac{1}{2}$ to $\frac{1}{9}$) with the addition of just one agent and commodity to the simulation.

It is arguable that more than mere chance is involved, even in the case of the random selection strategy. While the *selection* of actions for an agent is random, the point at which the selection is made is not – it is determined by how an agent is performing. Hence agents should still be able to react to each other's decisions and find a path towards mutually beneficial specialisation. The problem however, is that the agents all need to be increasing production of distinct commodities ('moving in the same direction') at the same time before the mutual benefits are seen. For a set of individual agents pursuing their own interests, this common alignment would seem to require an improbable and persistent coincidence of actions.

7.4.2 Need for a Heuristic

Given these findings, it would appear reasonable to apply one of the heuristic strategies suggested in section 4.2.2. If the only commodities that can be selected for increased production are those that an agent is selling, then specialisation happens fairly trivially as the following experiment shows:

Initial allocations

	Commodity 1	Commodity 2	Commodity 3	Specialisation Level
Agent 1	1.00	1.00	1.00	1.00
Agent 2	1.00	1.00	1.00	1.00
Agent 3	1.00	1.00	1.00	1.00

Results

Allocations after 700 production rounds:

	Commodity 1	Commodity 2	Commodity 3	Specialisation Level
Agent 1	3.00	0.00	0.00	3.00
Agent 2	0.00	0.00	3.00	3.00
Agent 3	0.00	3.00	0.00	3.00

Perfect specialisation is achieved and maintained as shown in figure 7.8.

This strategy promotes specialisation in distinct commodities, regardless of whether it is beneficial to the agents. In the case shown (with a production function that reflects the benefits of specialisation) specialisation *is* beneficial, to each agent individually and to the group as a whole. Given this assumption about production then, it would be quite reasonable for greedy agents to adopt such a strategy. In an evolutionary context, many adaptions are tailor-made to fit a particular environment and rely on built in assumptions about it – if the environment changes, the assumptions are challenged and the adaption may fail. In this case, the assumption is that specilisation is beneficial. If that proved not to be the case (for instance if generalising in different commodities resulted in higher productivity) then the above heuristic strategy would fail.

Evolving production strategies in a number of different such 'environments', would be interesting to see if similar heuristics emerge to cope with the combinatorial problem.

7.5 Summary

This chapter has presented results, analysis and discussion on a number of experiments involving production and specialisation. Two-agent specialisation was found to be fairly trivial, even with the basic production strategy. With three agents the coincedence in intent and actions required was too great and a heuristic strategy was proposed as an alternative.

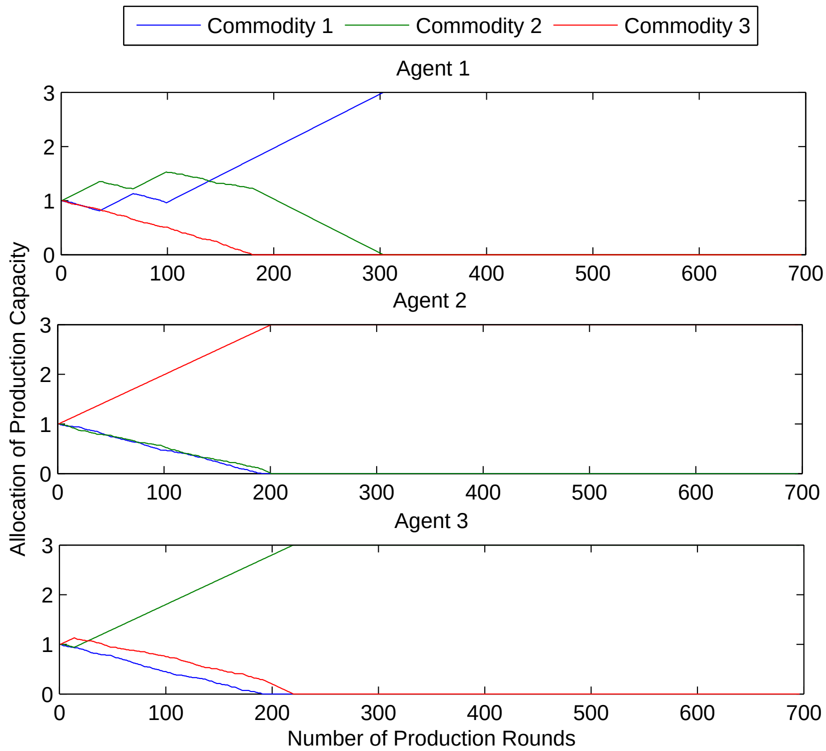

Figure 7.8: Allocations of three agents over 700 production rounds (Heuristic Strategy).

Chapter 8

Money

The emergence of money from networks of bartering agents was only investigated briefly in this project, but provided some promising points for further investigation.

8.1 Additions to the model

Some features were added to the model of production and specialisation, that might help to promote the emergence of money like commodities.

These included having different levels of **consumption**, allowing an agent to retain some of its commodities across production rounds (as opposed to the universal total consumption used previously). In particular, agents consumed the commodities they gained in trading, allowing them to keep from (round to round) whatever goods they produced and did not sell.

Different levels of **commodity decay** were also implemented, so that the commodities retained by agents decayed at some fixed rate per round. This stopped agents amassing huge stockpiles of a commodity and had the potential to determine whether different levels of decay led to different patterns of flow for a commodity.

To give some indication of how 'money-like' a commodity was, the percentage of transactions it was involved in is given. If this was close to 100 percent then almost every exchange of commodities in the network involved this commodity – one of the key properties of money.

8.2 Experiment

The following experiment implements the notion of consumption given above, with a ten percent decay rate across all commodities. It is conducted on a star network with one agent at the hub and three outlying agents connected to it (as shown in figure). All agents start non-specialised.

Initial allocations

	C1	C2	C3	C4	Specialisation Level
Agent 1	1.00	1.00	1.00	1.00	1.00
Agent 2	1.00	1.00	1.00	1.00	1.00
Agent 3	1.00	1.00	1.00	1.00	1.00
Agent 4	1.00	1.00	1.00	1.00	1.00

Agent 1 is at the hub of the star network, with tradelinks connecting it to Agents 2, 3 and 4.

Results

Allocations after 1000 production rounds:

	C1	C2	C3	C4	Specialisation Level
Agent 1	0.00	0.00	0.00	4.00	4.00
Agent 2	**2.23**	0.46	1.31	0.00	2.23
Agent 3	0.98	0.90	**2.12**	0.00	2.12
Agent 4	1.02	**2.16**	0.82	0.00	2.16

The percentage of transactions involving each commodity are

C1	C2	C3	C4
30.3%	33.3%	39.4%	97.0%

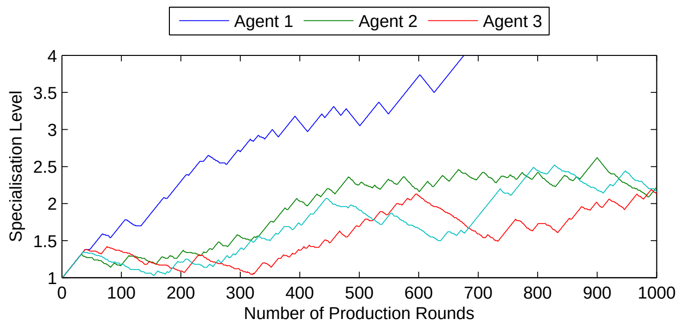

Figure 8.1: Specialisation levels of two agents over 1000 production rounds.

Analysis

The agent at the hub of the network specialised in one commodity, which became the dominant medium of exchange in the network (being involved in 97% of all transactions). The idea of a central, sole supplier of a commodity giving rise to money is an interesting concept, with obvious analogies to states and banking institutions. This occurs however, in a very limited set of circumstances. To see money arise more generally, I suspect some notion of resale value of a commodity would be required to promote its flow throughout athe network (as discussed briefly in chapter 3).

8.3 Summary

I added some features to the model intended to assist in the promoting the emergence and detection of money-like commodities. While few experiments were conducted, this is a fascinating area with a wealth of potential for future investigation.

Chapter 9

Summary

I have successfully defined and implemented a model of an agent-based bartering economy, and investigated the minimal factors needed to produce emergent behaviours and mutualisms within a network.

The overview of literature and projects revealed some interesting points to guide the investigation, but few took a similarly wide ranging approach to the problem as taken here. I defined a basic model of barter, in which simplicity and the *minimal* amount of features required, were priorities. This basic model allowed agents to trade commodities with each other in the context of a specified network topology. Production and specialisation aspects were added to the basic model definition, and a brief description of the implementation process given. Results of experiments on the model were presented and analysed, on the basis of statistics gathered.

9.1 Conclusions

Some important conclusions were drawn on the topics of 'the middleman', specialisation and money:

- The middleman exploits its position of power in the network to retain significant amounts of the commodities passing through the network. Its levels of profit were limited by the reluctance of the outlying agents to continue trading at such high rates of exchange.
- Competition between two agents in an *unrestricted* network topology results in a third party gaining a position of power (and increased profits). This happens because the third party can trade with *either* of the competing agents, who have little interest in trading with each other.
- Specialisation in 2 agent, 2 commodity simulations emerges straightforwardly, using even the simplest production strategy with the *least* assumptions.
- Specialisation in 3 agent, 3 commodity simulations does not emerge using this simple production strategy, as the combination of actions becomes too large for the coincidence in strategy required.
- If it is *assumed* that specialisation is potentially beneficial, then a heuristic strategy may reasonably be applied by a self-interested agent. Application of a particular heuristic strategy ('increase production of what you are selling') results in almost immediate specialisation throughout the network.

- Under limited circumstances (involving a centrally located, specialised 'distributor'), commodities may become a universal medium of exchange, one of the key features of money.

9.2 Future Work

There are a wealth of opportunities for future investigations on this topic. The project was quite open by nature and presented an array of possible options at each stage – both in the general approach and topics explored, as well as in the details of the model itself.

Following the basic approach outlined in this report, possibilities for future work include using negative commodity amounts, different utility and production functions, agents who do not necessarily fulfil their commitments, and different distributions of production capacity throughout the network. This is just a sample of the many possible options, outlined in more depth in Chapters 3 and 4.

Many of the additions to the model suggested had direct bearing on the issue of money, and the question of how money can arise from a basic system of barter is a wide ranging topic, with a huge number of possible paths to explore. Thes include lookin at decay and consumption rates, intrinsic properties versus resale value of commodities, the effect of different network topologies and notions of trust, quality and assymetric power differences.

Another main avenue to explore more fully is the effect that strategies (for both barter and production) have on emergent behaviour in a network. The interactions between strategies, and the evolution of strategies in a variety of simulated 'environments' would be of particular interest.

Bibliography

[1] AXELROD, R. *The evolution of cooperation*. Basic Books, New York, 1984.

[2] BANERJEE, A. V., AND MATKIN, E. S. A Walrasian theory of money and barter. *The Quarterly Journal of Economics 111*, 4 (November 1996), 955–1005.

[3] BONACICH, P. The evolution of exchange networks: A simulation study. *Journal of Social Structure 2*, 5 (2001).

[4] ENGINEER, M., AND SHI, S. Bargains, barter and money. *Review of Economic Dynamics 4*, 1 (January 2001), 188–209.

[5] FOSTER, J. From simplistic to complex systems in economics. *Cambridge Journal of Economics 29*, 6 (November 2005), 873–892.

[6] GARCIA-FLORES, R., AND KONTOLEON, N. Simulating a network of trading agents. In *5th Annual CSIRO CSS Workshop* (Sydney, August 2006). http://www.dar.csiro.au/css/Aug Workshop 06/'06 Workshop Presentations/ Rodolfo Garcia-Flores Trading Agents.ppt.

[7] GIBBS, S. Emergent mutualisms in an agent based economic simulation, 2005. BIT Project, Victoria University of Wellington.

[8] GRIMM, T., AND MITLOHNER, J. Developing a virtual reality for the simulation of agent based economic models. Tech. rep., Vienna University of Economics and Business Administration, 1995. http://mud.co.uk/dvw/developingavr.html (11 May 2006).

[9] JONES, P., AND BENSUSAN, H. OPTIMAES (open project to investigate money and economic systems). http://www.nooranch.com/synaesmedia/optimaes/optimaes.cgi (11 May 2006).

[10] LAVEZZI, A. Complex dynamics in a simple model of economic specialisation. Tech. rep., University of Pisa, May 2003. Discussion Paper no. 2.

[11] MCFADZEAN, D., AND TESFATSION, L. An agent-based computational model for the evolution of trade networks. In *Proceedings of the Sixth International Conference on Evolutionary Programming, Springer-Verlag, Berlin* (1997). http://www.econ.iastate.edu/tesfatsi/ep1997.ps (11 May 2006).

[12] PARKER, M. T. What is ascape and why should you care? *Journal of Artificila Societies and Social Simulation 4*, 1 (January 2001).

[13] PRENDERGAST, C., AND STOLE, L. Barter relationships. In *The Vanishing Rouble: Barter Networks and Non-Monetary Transactions in Post-Soviet Societies (Forthcoming)*, P. Seabright, Ed. Cambridge University Press, January 2000.

[14] RASMUSEN, E. *Games and Information: An Introduction to Game Theory*, 2nd ed. Blackwell, Cambridge, Mass., 1994.

[15] RUBENSTEIN, A. Perfect equilibrium in a bargaining model. *Econometrica 50*, 1 (January 1982), 97–110.

[16] RUBENSTEIN, A. A bargaining model with incomplete information about time preferences. *Econometrica 53*, 5 (September 1985), 1151–1172.

[17] TESFATSION, L. *Agent-Based Computational Economics: Growing Economies from the Bottom Up*. http://www.econ.iastate.edu/tesfatsi/ace.htm.

[18] WILLIAMSON, S., AND WRIGHT, R. Barter and monetary exchange under private information. *The American Economic Review 84*, 1 (March 1994), 104–123.