

UNIVERSITY OF THE WITWATERSRAND

DISSERTATION

SCHOOL OF COMPUTER SCIENCE AND APPLIED MATHEMATICS

**Emerging behaviours in Artificial Societies :
Simulating social-economic phenomena**

Author
B.C. Ingram

Supervisor
Prof. C Van Alten
Dr. R Klein

August 24, 2018



UNIVERSITY OF THE
WITWATERSRAND,
JOHANNESBURG

Abstract

The behaviour of a society and its individuals is inherently complex and, therefore it becomes difficult when trying to model how changes in a society will affect that population. This dissertation presents an implementation of a computer simulation technique known as an Artificial Society, which is used to model social-economic phenomena using a multi-agent system. An Artificial Society is a system of simulated agents in a simulated world governed by a set of rules which handle the nature of the interactions between entities of the system. The purpose of this model is to analyse the emergence of global behaviours that form from the evolution of the society as a result of interactions governed by simple rules. We firstly expand on a number of aspects of the work done by [Epstein and Axtell \[1996\]](#) such as trade and cultural transmission. We analyse the similarities of our model with that of pre-existing models such as the Predator-Prey model. We demonstrate the spread of culture through a society and investigate the influence an individual can have on a population in different environments. We also investigate and analyse the benefits and shortfalls that different kinds of trade, taxation and investment can have on a society in a 3 resource environment. Through all of these experiments, we were able to demonstrate the emergence of complex behaviours which formed as a result of the interactions of individuals. This approach of modelling whereby we generate simple rules and observe the emergence of global behaviour gives us the ability to tackle modelling complex behaviours where using a closed form solution would be impractical. It is hoped that this study will inform readers on the potentials and benefits of Artificial Societies.

Declaration

I declare that this Dissertation is my own, unaided work. It is being submitted for the Degree of Master of Science at the University of the Witwatersrand, Johannesburg. It has not been submitted before for any degree or examination at any other University.

Branden Corwin Ingram

BCI

August 24, 2018

Contents

Abstract	i
Declaration	ii
Table of Contents	iii
List of Figures	v
1 Introduction	2
2 Background	5
2.1 Artificial Societies	5
2.2 Cellular automaton	6
2.3 Game of Life	6
2.4 SUGARSCAPE model	7
3 Related Work	12
3.1 Early	12
3.2 Modern	15
4 Research Methodology	20
4.1 Methodology	21
4.1.1 Implementation	21
4.1.2 Testing	23
4.1.3 Analysis	24
4.2 Visualiser	24
4.2.1 Built-in visualiser	25
4.2.2 Standalone visualiser	26
5 The Simple Model	30
5.1 Introduction	30
5.2 Predator Prey Model Comparison	34
5.3 Survival of the Fittest	37
5.4 Supply and Demand	42
5.5 A comparison of movement	43
5.6 Birth Control	46
5.7 Conclusion	51
6 Culture	52
6.1 Introduction	52
6.2 The Emergence of a global Culture	54
6.3 Measuring the influence of individual agents	61
6.4 Cultural parameter as a model of information propagation	69

6.5	Cultural parameter as a model of information propagation based on topography	73
6.6	Conclusion	79
7	The Trade Model	81
7.1	Introduction	81
7.2	The benefits of Trade	85
7.3	Global trade versus local trade	88
7.4	Naivety	90
7.5	Trade Prices	91
7.6	Trading with preferences	94
7.7	Trading based on ideology	96
7.8	Conclusion	99
8	Tax	100
8.1	Introduction	100
8.2	Wealth distribution control study	102
8.3	Taxation versus Investment	104
8.4	Different Tax rates	106
8.5	Progressive Tax Rates	108
8.6	Indirect Wealth Redistribution	118
8.7	Conclusion	120
9	Conclusion	121
	References	123

List of Figures

2.1	Basic combinations of live cells when applying the rules of Game of Life Gardner [1970]	7
2.2	Initial state of the SUGARSCAPE acting on the movement and regrowth rules	9
2.3	SUGARSCAPE after a couple of episodes with the movement and regrowth rules	10
2.4	SUGARSCAPE after a couple more episodes with the movement and regrowth rules	11
3.1	R-pentomino shape, the dark cells represent the live cells [Martin 2013]	12
3.2	R-pentomino shape, after 20 generations [Martin 2013]	13
3.3	R-pentomino shape, after 410 generations [Martin 2013]	13
3.4	R-pentomino shape, after 1103 generations [Martin 2013]	14
3.5	R-pentomino shape, after 1216 generations [Martin 2013]	14
3.6	Migration of agents on the SUGARSCAPE [Epstein and Axtell 1996]	16
4.1	Built-in visualiser	25
4.2	Built-in visualiser	26
4.3	Standalone visualiser menu	26
4.4	visualisation of sugar grid without agents	27
4.5	Visualisation of agent's sugar values	28
4.6	First visualisation of agent's spice metabolism	28
4.7	Visualisation of agent's gender, blue = male, red = female	29
5.1	Visualisation of the replica of the SUGARSCAPE model depicting the levels of sugar	33
5.2	Predator Prey model	35
5.3	Predator Prey model realistic conditions	36
5.4	Predators (population of agents) and Prey (amount of sugar) comparison	37
5.5	Number of agents grouped by metabolism for the simple model	38
5.6	Average wealth of agents grouped by metabolism for the simple model	39
5.7	Number of agents grouped by vision for the simple model	40
5.8	Average wealth of agents grouped by vision for the simple model	40
5.9	Number of agents grouped by metabolism for the simple model with the more realistic mutation	41
5.10	Total sugar available and the total metabolic demand comparison for the simple model	42
5.11	Total number of agents for each movement rule	44
5.12	Average age of agents for each movement rule	44
5.13	Average wealth of agents for each movement rule	45
5.14	Total amount of sugar for each movement rule	46
5.15	Total number of agents per the differing birth control schemes	47
5.16	Average wealth of agents per the differing birth control schemes	48
5.17	Change in population size with introduction of 2 child policy at episode 500	49
5.18	Total Births and Deaths with introduction of 2 child policy at episode 500	49

5.19	Average age and wealth with introduction of 2 child policy at episode 500	50
6.1	Total number of cultural differences for the simple model	54
6.2	Total number of cultural differences where mutation was prevented for the simple model	55
6.3	Visualisation of the Simple Model using the Influence Rule, Episode 0, 5, 15 and 27	56
6.4	Visualisation of the Simple Model using the Influence Rule, Episode 35, 42, 105 and 130	58
6.5	Visualisation of the Simple Model using the Influence Rule, Episode 135, 416 and 996	60
6.6	10% chance of Information propagation with “special” agent	63
6.7	0.1% chance of Information propagation with “special” agent	65
6.8	0.1% chance of Information propagation with “special” agent and vision 10	67
6.9	10% chance of Information propagation	70
6.10	0.1% chance for Information propagation	72
6.11	Cost of spreading information in the long bridge world	74
6.12	Long Bridge, 10% chance experiment	75
6.13	Long Bridge, 1% chance experiment.	77
6.14	Long Bridge, 0.1% chance experiment.	78
7.1	Levels of sugar across the world. Darker yellows indicate higher yields	82
7.2	Levels of spice across the world. Darker yellows indicate higher yields	82
7.3	Levels of food across the world. Darker yellows indicate higher yields	83
7.4	Total population of trade vs no trade for Trade Model	86
7.5	Average age of trade vs no trade for Trade Model	87
7.6	Average wealth of trade vs no trade for Trade Model	87
7.7	Total population per trading strategies for Trade Model	88
7.8	Average age per different trading strategies for Trade Model	89
7.9	Average wealth per trading strategies for Trade Model	89
7.10	Average Intelligence for Trade Model	91
7.11	Average trade prices for Trade Model	92
7.12	Average trade prices with dynamic events for Trade Model	93
7.13	Number of cultural differences grouped by experiment	95
7.14	Number of trades grouped by experiment	95
7.15	Population grouped by experiment	96
7.16	Number of agents divided by their trading strategy	98
8.1	Initial wealth distributions	102
8.2	Wealth distributions of control experiments	103
8.3	Wealth distributions of differing tax strategies	105
8.4	Population for different wealth redistribution strategies	106
8.5	Wealth distribution for different tax rates	107
8.6	Population for different tax percentages	108
8.7	Population for tax plan 1	110
8.8	Population for tax plans where tax was introduced after episode 200	110
8.9	Total Births for tax plans where tax was introduced after episode 200	111
8.10	Population for tax plans with low initial endowment	112
8.11	Wealth distributions of Tax plans with low initial endowment	113
8.12	Population for tax plans with low initial endowment and taxation only after episode 200 .	114
8.13	Wealth distributions of Tax plans with low initial endowment and taxation only after episode 200	115
8.14	Population for tax plans with fixed tax brackets	116
8.15	Wealth distributions of Tax plans with fixed tax brackets	117

8.16 Population and wealth distribution for Indirect Wealth Redistribution experiment	119
8.17 Wealth distribution at Episode 1000 under environment upliftment	119

Chapter 1

Introduction

[Epstein and Axtell \[1996\]](#) were the first to use agent-based systems for modelling social sciences. They wished to investigate how individuals' behaviour influenced society as a whole. That is, if individuals or agents follow specific rules, what properties will the society ultimately develop. They introduced "Artificial Societies" which are an agent-based computer modelling technique used to study in particular social phenomena such as trade, the spread of culture and information, evolution, combat and disease transmission. Artificial Societies may be considered a multi-agent system consisting of an environment together with a collection of agents that interact with each other and the environment based on a fixed set of rules. Such a multi-agent system is well suited for implementation on a computer as the power of modern computers has made this approach more feasible.

One aspect of an Society is that individuals within the society are not all the same and they do not all interact in the same way as others. To make modelling realistic for a real-world scenario we want to model individuals who are not perfectly defined nor should they be perfectly informed. The flexibility of multi-agent systems allows us to incorporate these factors by defining agents and an environment with both internal parameters and rules, all of which can be adjusted to fit the simulation of our Artificial Society. This system can easily be implemented using Object Oriented Programming as this technique fits naturally with the multi-agent based models. Additionally, various aspects of a society of people can be modelled in conjunction and need not be decomposed into separate studies, thus leading to more realistic models. This allows for complex and dynamic social processes to be modelled and analysed together, from which we hope to gain new insights into social sciences.

Modelling with Artificial Societies can be used as a prediction model to help us better understand the effect of the introduction of something new in a society or the change in pre-existing aspects. For example, the government could be looking to implement a new welfare bill whose effect on the entirety of a country would be very hard to predict. However, if we generate an artificial society that operates without this bill, record some data, then run the same simulation with this new bill included, we can then get a better understanding of the effect by comparing the new data recorded to that of the previous simulation. Due to the flexibility of the model, we can also tweak rules and parameters of the society as well as make adjustments to the bill to possibly correct shortcomings discovered by the simulation. The second reason becomes one of ethics. We cannot simply remove a food supply from a community of people and watch to see the results. Nor can we induce a natural disaster onto a society and wait for those results. Not only would it be inhumane but time constraints would also be a substantial problem. One cannot wait years for emergent behaviour to arise and use that as an accurate model, especially in real life as other unknown dynamics might have come into play during that period. The first and second reason revolved around analysing emergent behaviour to use for future predictions. However, there is

another application: that of using this model to make predictions of the past. For example, one could attempt to model the way primitive man interacted with his environment using what knowledge we have accrued in order to generate a model that can be used to monitor the migration of mankind in its earliest infancy. So instead of using it as a form of future prediction, we can use it as a method for history generation.

This research demonstrates the efficacy of using this approach to model the emergence of complex socio-economic behaviours through the interaction of agents. These include dynamics such as welfare, trade, the spread of cultural beliefs, effect of differing changing environments, population control and class structures in terms of both economic and cultural affiliation, all of which are discussed in later chapters.

Another heavily debated problem in modern society that we analysed was that of overcrowding. This problem was extended from just overcrowding to include the ever increasing demand on the world's natural resources. For this problem, we looked to the dangers of over-consumption and how that might lead to poverty or even extinction level events. Secondly, we analysed the effects of enforcing different birth control schemes on populations much in the same way as China's one-child policy. These experiments were conducted to see how the level of population would be affected as well as their quality of life. Finally, we considered the ethical nature of the implementation of population control mechanics. We observed that population control measures had a big impact on population sizes as well as average wealth.

Cultural transmission was simulated by allowing agents to affect a cultural parameter of other agents, thereby simulating the effect other individuals have on one's own belief system. For cultural transmission, we found that groups of agents form with similar cultural values and over a longer period of time a single uniform belief across the entirety of the society formed. We also determined how an individual's beliefs can have a large impact on the global belief system. For example, we determined who might be the most influential in a society. Tracking influence in a society is an important problem, especially in today's landscape where people are bombarded with a large number of influences from social media. This problem led to additional simulations, which involved analysing the propagation of information through a society. The society we live in is a connected network in which we are the nodes and the mediums of communication form our edges. With this idea, we felt our multi-agent system was ideal for modelling this propagation of information in different types of environments. Here we analysed two aspects; how do individuals affect how information is spread as well as how does the landscape they live in affect it. We observed that individuals have the potential of influencing the global belief system of the society how landscape has a notable impact on the spread of information.

Trade was handled by allowing agents to trade goods among other neighbouring agents or perhaps another entity representative of a shop. What we witnessed here is that real-world market trends and dynamics arose such as the formation of a global market price for differing goods. Another observation was the changing of average wealth of those agents who partake in trade and those who refuse. What we concluded from here was that of allowing agents to trade increased their average wealth. Through our experimentation with trade we concluded that allowing trade resulted in a positive benefit to the society.

Finally, the last aspect was that of the impact of taxation. Taxation simulations are run to monitor the effect of introducing welfare initiatives on certain members of the society. For example, the old or poor might receive a small amount of wealth. In this case, we might hope to see that if we supplement the wealth of the poor and old, the average age or population size of the society might increase. On the other hand, we might also note changes in the average wealth by critically looking at distributions. We also compare the notion of higher taxation on the rich with the contending belief that this form of taxation in fact negatively affects the poor. Experimentation with taxation led us to conclude that lower tax rates were beneficial but also that taxation increased population sizes to a degree.

Multiple simulations are run on a variety of environments ranging from ones with scarce resources to others with an abundance of resources. Not only can resource levels be simulated, but we can also simulate differing terrains and landscapes where movement becomes dependant on the terrain in question. Here a possible conclusion would be that carrying capacity of a society is greatly affected by the level of abundance of resources. In environments with scarce resources, only a small amount of agents are sustainable and, in even worse cases, extinction events may occur. The opposite is predicted for environments of abundant resources.

Chapter 2 introduces the background knowledge and concepts relating to the previous work done within the field. Following this introduction, Chapter 3 describes how Artificial Societies have already been used in order to model the different effects of various socio-economic behaviours of individuals and populations. Such socio-economic behaviours include trade amongst different entities, the emergence of cooperation, the effect of a territorial behaviour on tribes of agents, taxation and application of Artificial Societies in Prisoners Dilemma. Chapter 4 describes in depth the research plan required to test our experiments. This includes the different questions we looked to analyse as well as giving a description of the design of our model. Finally, it outlines our approach to generating and storing data which can be used to achieve our goal of observing emergent behaviour. The last 4 chapters all revolve around investigating problems and studying the emergence of interesting behaviours relating to the main topic specific to that chapter. The first is Chapter 5, where we looked into replicating the SUGARSCAPE as well as comparing our model with other existing models. We also looked at how our model could be utilised to study policies which affect population dynamics and structures. Chapter 6 revolves around studying the effect of culture in a society and how it can be used to model information propagation. Chapter 7 is about how allowing agents the ability to trade leads to the emergence of a market like structure. Our last chapter of results being Chapter 8 looks to analyse the effect of taxation versus investment in a society. Here we took a deeper look into the success and failure of a society at an individual level. Chapter 9 concludes the document by summarising the interesting behaviour observed together with touching upon some limitations and future work.

Chapter 2

Background

This chapter introduces the different concepts relating to an Artificial Society. Examples are given to better explain these concepts and show how they can be used to develop Artificial Societies that exhibit sensible emergent behaviour on a global level.

2.1 Artificial Societies

An Artificial Society is a multi-agent based computer modelling technique used to study social behaviour. Social behaviour in essence is the interaction between entities, which can range from people all the way to more abstract ideas like companies or governments. These entities are represented by the “agent” component of a multi-agent system. A multi-agent system is a computerised system which consists of three fundamental elements: agents, their environment and a set of rules. These rules govern the way the environment and agents interact. The origin and development of Artificial Societies will be expanded upon in Section 2.3 and 2.4. In Chapter 5 we describe in detail the design of a simple Artificial Society using a multi-agent system. [Epstein and Axtell 1996]

An agent is an object which consists of an internal set of parameters which form the agent’s attributes. These attributes are common across all agents, however, the values for attribute may vary. For example, since an agent represents real world entities such as people, sex, metabolic rates as well as vision could all be considered as part of a person’s internal parameters or attributes. These parameters may change over time dependant on what they represent. The initial set of agents form the initial population of the model.

Typically an agent will exist at a particular cell position within an environment at any given moment in time. The environment is also an object which is made up of a network of cells and is typically represented by a 2D grid where each grid coordinate is a cell. These cells all have their own internal parameters much like that of the agents. However, the environment also has a set of global parameters which each cell shares. A cell’s internal parameters contribute to defining that particular cell whereas the global parameters contribute to defining all of the cells.

The final component of a multi-agent system is the set of rules. Rules are the instructions which govern the interaction between agents and other agents, between agents and the environment as well as between the environment and itself. Essentially a rule defines an action to be made based upon the internal parameters of the two interacting objects.

This means that the decisions that agents make (an agent’s behaviour) are based upon all their internal

parameters as well as being dependant on the current state of the environment and other agents. Additionally, this means that the current state of the environment is a result of all the interactions with agents which have already occurred as well as it's own internal parameters. The way in which the next state is determined is based on the current state of the environment, the set of agents as well as the rules.

2.2 Cellular automaton

A cellular automaton is a discrete model consisting of a regular grid of cells, each in one of a finite number of states, such as on and off. The grid can be in any finite number of dimensions. A set of cells defined with respect to a specified cell is called its neighbourhood. An initial state at time $t = 0$ is selected by assigning a state for each cell. A new generation is created by advancing t by 1 according to some fixed rule or set of rules. These rules determine the new state of every cell in terms of the current state of the cell as well as its neighbourhood. Typically, the rule for updating a cell is the same for each cell, does not change over time, and is applied to the whole grid simultaneously, however, this may not always be the case. Understanding of cellular automaton is important because an Artificial Society is essentially a complex form of cellular automaton.[Gardner 1970]

2.3 Game of Life

Game of Life can be considered more of a simulation than a game because it is a fully automated zero-player system. It is still, however, a great example of a simple multi-agent system which demonstrates how the three elements discussed above work together. The only user input required is the initial state of the world. From that point, the evolution of the system is fully autonomous, requiring no further input.

To “play” with the Game of Life a player creates an initial configuration and then observes how the state of the world evolves over time. The environment of the Game of Life is an infinite two-dimensional orthogonal grid of square cells. These cells can each be in one of two possible states. An analogy for this that each cell can either be dead or alive. Every cell interacts with its eight neighbours, which are the cells that are horizontally, vertically, or diagonally adjacent. This 2D grid is governed by simple rules which are mentioned below. At each step in time, the following transitions occur:

- Any live cell with fewer than two live neighbours dies, as if caused by under-population.
- Any live cell with two or three live neighbours lives on to the next generation.
- Any live cell with more than three live neighbours dies, as if by over-population.
- Any dead cell with exactly three live neighbours becomes a live cell, as if by reproduction.

The initial pattern of the environment constitutes the seed of the system. The first generation is therefore all cells that are alive on the seed system. The next generation is calculated by applying the rules described above to all cells simultaneously. The rules continue to be applied repeatedly to create further generations. This means that each generation is a function of the previous. Consider Figure 2.1 below, more specifically the grid a_0 . By applying the rules to the top-most live cell, we note that it has fewer than two live neighbours therefore it dies. This conclusion can be seen in the grid a_1 , where the same cell is no longer shaded.

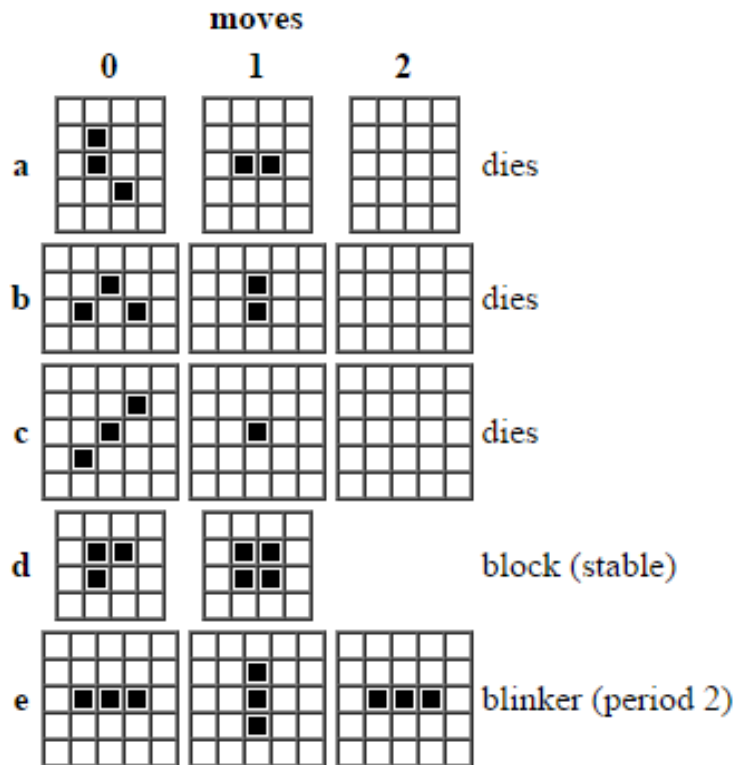


Figure 2.1: Basic combinations of live cells when applying the rules of Game of Life [Gardner \[1970\]](#)

2.4 SUGARSCAPE model

The second example of a multi-agent system is the one of [Epstein and Axtell \[1996\]](#) who pioneered the idea of an Artificial Society. They utilised multi-agent systems and object-oriented programming in their attempt to model a variety of socio-economic phenomena such as reproduction, the spread of culture, trade dynamics, conflict among agents, the gathering of wealth, the spread of disease and the generation of history. They dubbed this model the SUGARSCAPE, aptly named after the primary resource required by the agents in order to survive. In their implementation, they represented people by agents and the environment they lived on as 2D grid lattice on which the agents could interact. In essence, agents were able to move and collect sugar placed out in different yields on the grid in order to survive. These agents also had other rules which governed their interaction as well as internal parameters which affected these rules in a number of ways. Descriptions of the internal parameters and rules are provided below.

Parameters

- Sugar
- Metabolism
- Vision
- X, Y (represents an agent's location)
- Culture (an array of values)
- Gender

Agent based Rules The agents of the SUGARSCAPE acted under a number of rules. How these rules functioned is described below.

The movement rule worked as follows :

- Step 1: The agent would look out on the grid as far as its own vision parameter permitted in all of the four primary movement directions (up, down, left and right).
- Step 2: Identify the vacant tile with the highest yield of sugar.
- Step 3: Move directly to that tile by adjusting the agents x and y coordinate.
- Step 4: Collect all the sugar at that new tile.

The next rule they implemented was the reproduction rule. This rule worked as follows:

- Step 1: Agent would identify if there are any neighbouring agents.
- Step 2: For each neighbouring agent, a check would be done to see if reproduction could occur based on parameters like gender, age and wealth.
- Step 3: If this check passed a new agent would be born with parameters based on that of his/her parents and placed on any vacant tile neighbouring their parents.

The rule which governed cultural transmission worked as follows:

- Step 1: Agent would identify if there are any neighbouring agents.
- Step 2: For each neighbouring agent an element in the culture array is randomly selected.
- Step 3: If the neighbour agrees with agent at that position, in other words, they have the same value at that position, no change is made. If they disagree, in other words, have differing values, the neighbour's value is changed to match that of the current agent.

The rule which governed combat worked as follows:

- Step 1: Agent would look out on the grid as far as its own vision parameter permitted in all of the four primary movement directions (up, down, left and right).
- Step 2: Ignore all sites (2D grid locations) occupied by members of the agent's own tribe.
- Step 3: Ignore all sites occupied by members of different tribes that have higher wealth values. An agent's tribe was determined by the culture parameter. This was to simulate the idea that you would not attack someone relatively stronger than yourself.
- Step 4: The reward of each remaining site is given as the amount of sugar at that site plus half the sugar amount of the agent, if any, occupying the tile.
- Step 5: Move to the tile with the highest reward, if the tile is occupied the occupying agent would be killed and the agent would receive half the dead agent's sugar.

Environment based Rules

The rule which governed how the sugar grew-back worked as follows: At each grid position the sugar level grew back at a rate of α units per interval up to the capacity of the grid position.

The Bigger Picture

Further understanding of how all these aspects interacting together actually work is obtained from Figures 2.2, 2.3 and 2.4 below. A description of what is happening and why follows each figure as well as

how the rules and parameters described above are affecting the interaction. The initial sugar yields can also be seen which is represented by varying shades of yellow. The highest yields are the brightest shade of yellow situated close to the top right and bottom left regions of the environment.




Figure 2.2: Initial state of the SUGARSCAPE acting on the movement and regrowth rules

Figure 2.2 above represents the initial state of the model. The agents have been initially placed on the grid in random unoccupied locations. This point is seen by the distribution of the agents across the entire environment. Once this initialisation phase has been completed the simulation can begin. The cycle of each episode goes as follows: for each agent we first consume an amount of sugar based on that agent's metabolism parameter, then we execute the movement rule. The environment then executes the regrowth rule. This cycle continues a given number of times.

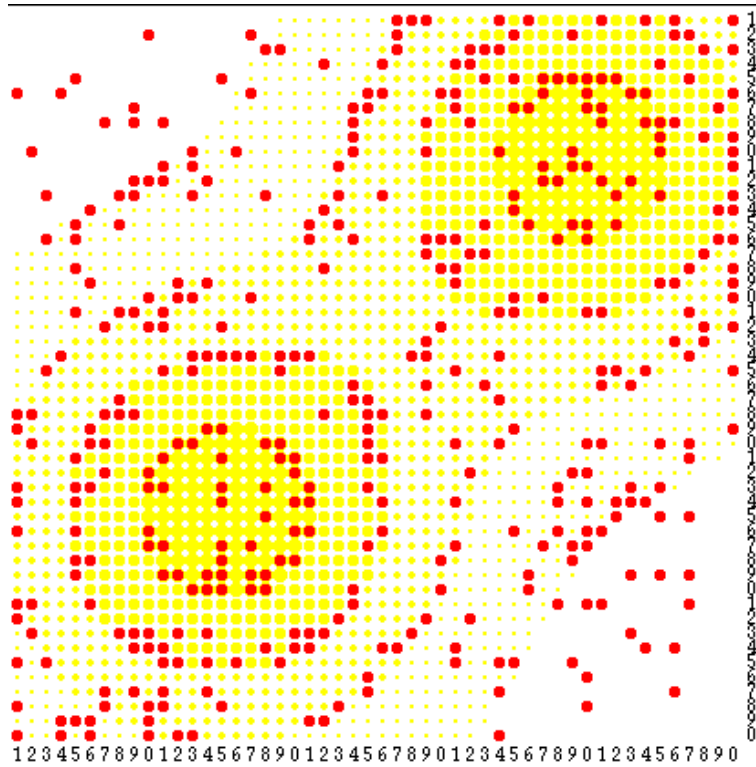


Figure 2.3: SUGARSCAPE after a couple of episodes with the movement and regrowth rules

In Figure 2.3, the effects of the movement rule can be plainly seen as agents have already begun to congregate towards higher yielding tiles. We also note that some agents remain stationary on the edges of the terraces. This is due to the vision parameter as they have a value which is low enough to prevent them from seeing the higher yield region as stipulated by the movement rule. We also note that some agents don't move away from the 0 yielding regions, which is also due to the vision parameter preventing them from seeing any better options than that of their current location. A final observation can be made that the population decreases, which is due to agents dying from starvation or running out of sugar, the reason not all agents in the 0 yield regions die at the same time is due the metabolism parameter. Agents with a low metabolism will consume less sugar per episode and therefore survive longer.




Figure 2.4: SUGARSCAPE after a couple more episodes with the movement and regrowth rules

When looking at Figure 2.4, we see that only a few agents have survived when compared to our initial population. This is due to the fact that those agents who die were never replaced either directly or through the reproduction rule. There is also a higher percentage of agents on the higher yielding regions, which is because these regions can sustain agents with a high value for their metabolism parameter. The movement rule also promotes the migration to these regions. We can also see that the yields have changed slightly whereby there are far fewer of the max yield regions than before. This is due to agents consuming the sugar at these regions, along with the regrowth rule. Agents are consuming sugar faster than the environment is regrowing it which has resulted in these white 0 yield tiles appearing in place of the higher yielding tiles. For this simulation an α of 1 was used. A final important observation is that of the surviving agents on the outskirts of the environment in the low yield regions. This is the result of these agents having a low vision as well as a low metabolism, meaning these agents can survive in the low yield regions indefinitely however they can't escape them either as they can't see a better alternative.

Chapter 3

Related Work

This section outlays a brief history of Artificial Societies as well as some of the more notable studies that have been done within the field. The studies discussed below were chosen as they share a similar design to that followed in this work. The early work focuses more on cellular automata and particle-based systems, whereas, the more modern approaches utilised complex multi-agent systems in order to model complex social behaviour.

3.1 Early

The earliest forms of agent-based models can be traced back to Von Neumann machines [Von Neumann and Oxtoby 1988]. He proposed a device which would follow precise rules in order to make a copy of itself. Von Neumann’s hypothetical machine used a “sea” of spare parts as a source of raw materials. The machine then was directed by a program stored on a memory tape. This program would hypothetically direct the machine to retrieve parts from the “sea” using a manipulator and then assemble them into a duplicate of itself. Finally, it would then copy over the program to the memory tape of the new duplicate machine. This thought experiment only became popular as a result of Penrose [1959]. A mathematician named Stanislaw Ulam (a friend of Von Neumann) suggested building his machine on a 2D grid. Von Neumann took this idea and thus created what would go on to be called cellular automata. From this, we learned that seemingly complex environments can still be represented on a 2D grid.

A further advance was achieved by mathematician John Conway who developed the well-known game of life in 1970 [Gardner 1970]. Conway’s Game of Life, as discussed in Chapter 2, operated in the form of a 2D grid which was governed by very simple rules. He chose rules that produced the most unpredictable behaviour. One of the most interesting early shapes was the R-pentomino.




Figure 3.1: R-pentomino shape, the dark cells represent the live cells [Martin 2013]

This pattern was discovered when testing all the different combinations of patterns with less than 6 live cells. It was by far the most active combination taking 1103 generations to stabilise, whereas all other combinations would stabilise in at most 10 generations. By “active” we are referring to when the pattern of live cells changes from generation to generation and a combination is referred to as “stable” when it no longer has any change in live and dead cells. Figures 3.2,3.3,3.4 and 3.5 below demonstrate the evolution of the R-pentomino shape over time as well as the formation of stability after 1103 generations.




Figure 3.2: R-pentomino shape, after 20 generations [Martin 2013]




Figure 3.3: R-pentomino shape, after 410 generations [Martin 2013]




Figure 3.4: R-pentomino shape, after 1103 generations [Martin 2013]




Figure 3.5: R-pentomino shape, after 1216 generations [Martin 2013]

From Gardner [1970] we gained an understanding of how interactions on a simple 2D grid can lead to the emergence of interesting behaviours. For example in Figure 2.1 we can see how some initial patterns can lead to extinction events namely rows a, b and c. The pattern depicted on row d also lead to an interesting behaviour whereby it remains constant and stable after a single generation. The last row depicts a different stable behaviour whereby the population oscillates forming a blinking pattern. Greatly complex behaviours are also present in the simple Game of Life framework such as those seen by the R-pentomino pattern. Throughout the evolution of R-pentomino, we see the migration of live cells until ultimately a stable pattern forms. All these behaviours emerge from the simple interactions of live and dead cells based on the rules previously described.

Craig Reynolds, a computer scientist, was the initial driving force behind the application of the agent-based model as a social model [Reynolds 1987]. He attempted to model living biological agents, a

method known as artificial life. He was able to generate a system which modelled flock, herd and school dynamics of animals using particle systems. This model allowed him to accurately predict the movement of simulated flocks of birds. This was achieved by implementing the birds as independent actors or agents that determined their own path based on a rule. He found that this approach was more realistic than individually scripting the paths of the birds. This research led to further studies on swarm mechanics using both particle-based systems as well as agent-based models.

[Epstein and Axtell \[1996\]](#) applied the computational methods of artificial life to the analysis of social systems in their book “Growing Artificial Societies: Social Science From the Bottom Up”. This book will be discussed in more detail below.

3.2 Modern

The artificial society has been widely accepted by recent sociology as a promising method characterised by the extensive use of computer programs and computer simulations which include evolutionary algorithms, genetic algorithms, genetic programming, memetic programming, agent-based models and cellular automata [[Epstein and Axtell 1996](#)].

This, however, was not always the case and, as mentioned above, the book by [Epstein and Axtell \[1996\]](#) can be considered the first research done in Artificial Societies. The aim of their book was to begin the development of studying human behaviour from an evolutionary point of view by applying a computational approach. The multi-agent system they designed as discussed in [Chapter 2](#) utilised advanced rules when compared to [Gardner \[1970\]](#). These advanced rules were used for a number of differing simulations, however, they remained simple in design and easy to implement. They were successful in their research as they were able to simulate many real-world phenomena through the emergent structure and behaviour that evolved in their Artificial Society.

An example of such emergent behaviour was population dynamics, such as the carrying capacity of a world, which would tend to some constant or would form a stable oscillation. A unique simulation was run where all agents were placed together in the bottom left region of the grid, while only considering the movement, reproduction and regrowth rules.




Figure 3.6: Migration of agents on the SUGARSCAPE [Epstein and Axtell 1996]

Here they concluded that wave-fronts of agents would move out periodically dependant on the regrowth rate of the world. These wave-fronts were considered to emulate the migratory patterns of nomadic man. They also noted that the high vision agents would migrate while the low vision agents would not. They proposed this behaviour to match that of the animals, where birds with their higher vision migrate whereas bears with a lower vision hibernate.

In another simulation which considered the culture rule, a form of tribalism was the emergent structure that evolved. This was seen by colouring agents with similar cultures with the same colour. What evolved, in this case, was that agents which operated in similar regions of that of others would have the same colour. In some cases the entire world would converge to a single colour, meaning all agents shared the same beliefs or culture. Here we noted that individuals have an equal chance influencing each other, which is unrealistic for a real world and as a result of this we looked to model people who have greater and lesser influence on others.

Market dynamics such as market price equilibriums evolved in a two resource society by running a simulation which took into account the trade rule. They found that over the course of the simulation the trade price of each commodity would tend towards a global market price. From the experiments done involving trade we wanted to see if their results could be obtained if we expanded their model to include a third resource.

Biological evolution was a staple ingredient in all the simulations. This was due to the fact that when an agent was born their birth parameters such as vision and metabolism were dependant on their parents. This meant that particular biological traits that were successful would survive while worse traits like low vision died off. In the end, the society that would evolve would only contain agents with very similar values for their parameters. Here we noted that the evolution observed is not perfectly consistent with Darwinian Evolution as the population starts off very diverse. Based upon this we looked how we could generate a perfect match for Darwinian Evolution as well as how Artificial Societies as a model can compare to other models.

From the research of [Epstein and Axtell \[1996\]](#) has come thousands of subsequent papers as well as a whole journal dedicated to Artificial Societies. The Journal of Artificial Societies and Social Simulation (JASSS) is a quarterly peer-reviewed academic journal created by Nigel Gilbert in 1998. According to ISI Journal Citation Reports, the impact factor of JASSS in 2017 was 1.640 [[Squazzoni 2018](#)]. Their work serves as the foundation of this research and has already had a major impact in the field having been cited 5096 times as per Google Scholar.

[Rouchier \[2002\]](#) presented a paper called “Multi-agent systems for the study of social complex systems”. The aim of her research was to analyse the different patterns of relations that could be created simply through merchant exchanges. These exchanges would occur in an artificial world that captures the situation of herdsmen trying to have access to a resource over time. In essence, the goal was to simulate the emergence of relationships through trade history. She implemented a world which contained three populations, one mobile population (herdsmen) and two offering services (village leaders, farmers). The goal of the herdsmen was to grow their flock. This was considered to be “production” and this was achieved by interacting with the other members of the world. This interaction was required in order to get water from the village leaders and land to graze from the farmers. The rational abilities of the agents were restricted so that the only kind of agent interaction is a dialogue between an artificial herdsman who proposed a deal to an artificial farmer, who answered. To build the agents’ logic, two ways of reasoning were implemented for their agents: one inspired by the idea of cost and the other by the idea of relation. To capture the idea of costs, any refusal is integrated by the herdsman as an extra-cost anticipated for the new transaction. To capture the idea of the value of a relation, the herdsman remembers the refusals and the agreements made by the farmers.

Through her research, she was able to show that genuine market trends like competition arose within the system as certain herdsmen were vying for particular farmers. She also discovered that competition had a negative effect on production. She proposed a possible reason for this, being the way in which the memory of refusals effects cooperation. In her implementation, she notes that a herdsman forgets all the refusals that a farmer made just after another farmer accepts one proposition. She concluded this inability to remember the accumulation of refusals is the reason for a herdsman’s memory being less adaptive as when the agent gets refused its point of view changes too quickly.

[Buzing et al. \[2005\]](#) investigated the evolution of communication between agents under various levels of “cooperation pressure” in an artificial society. In other words, they wanted to test the varying successes or failings of different societies in which they controlled the level to which agents were able to communicate with each other. They utilised the standard SUGARSCAPE model as described above. In their study agents were given the possibility to “call for assistance” if they encountered a sugar pile too large

to be eaten alone. This possibility was determined by a “talkativeness gene” which determined to what extent a given agent is willing to talk to others. The same concept was applied to a “listening gene” on the other end of the communication channel. A high cooperation pressure meant agents were given high values for their listening and talkativeness gene, this would promote cooperation. The opposite can be said for low values. Now by implementing agents which are given a random value for both genes, they can see which values for these parameters perform best over a given time.

They were able to determine that societies that have the ability to communicate (to talk and to listen) are viable in a significantly wider range of circumstances. They also found that only through communication were agents able to survive in environments with increasingly scarce resources. They also saw that the willingness to listen emerges earlier than the willingness to talk. This means that a society of people who are more willing to listen arose sooner than that of a society who is willing to talk. Finally, they compared their findings to that of the previous non-multi-agent system approaches namely [Oliphant \[1996 1997\]](#) and were able to see correlations in their results. That being in both approaches a willingness to communicate emerges. They noted that send and receive functions that Oliphant presents were very similar to the probability functions that we used in the controllers of the agents for talking and listening. However, they noted that their talking and listening functions were somewhat simpler.

A further extension of the SUGARSCAPE model was done by [Klöck \[2007\]](#) in his paper called “Extended sugarscape model in XL”. His aim was to study how territorial behaviour effects cultural transmission, combat between individuals in different groups and trade. He decided to run multiple simulations each focusing on comparing the effect of one aspect on another e.g. territorial behaviour versus trade. He decided to implement everything using the programming language XL a rule-based extension of Java in the GroIMP environment. The GroIMP environment is a 3D modelling platform written in Java which is useful for the algorithmic modelling of 3D structures. He states that the biggest difficulty was the usage of the language as it lacked many tools and libraries available in many other more common languages such as a debugger. He was able to determine that when he was comparing territorial behaviour versus cultural transmission, the amount of transmission was dependant on the degree to which agents of the same group were separated. He found when agents of each group were randomly spaced across the world that cultural transmission was much easier and was not affected by the territorial behaviour. This, however, changes when all agents were initially placed together. He was able to see a decrease of cultural transmission in this case. Following this, he compared combat versus territorial behaviour and determined by implementing the territorial behaviour that far fewer conflicts arose. He deduced that this was because tribes only fight if their relative power shifts and even if this occurs only a few regions would be invaded. The relative power of a tribe being their total wealth, so if an opposing tribe’s power decreased it would be more likely to be invaded by a relatively stronger tribe. Finally, he, compared the effect of adding trade into the same simulation. In this simulation, each tribe had its own resource, which meant that one tribe starts with the sugar field, and the other with the spice. This meant that the agents would have to trade or fight with the other tribe in order to survive. In order to incorporate the territorial aspect one of the tribes was able to own regions. At first, there were only intra-tribe trades occurring, then at some point when the tribes met for a period of time there were trades between tribes, however, this did not last long. This is because the tribe that cannot buy land dies off, which is a result of not being able to reach the sugar due to the territorial ownership of the other tribe as well. From this simulation he concluded that trade was not an important aspect, the aspect that was more important was the fact that regions prevented opposing tribes from accessing resources.

A book called Think Complexity by A. B. Downey applied the SUGARSCAPE model to another economic phenomenon, namely taxation, [[Downey 2012](#)]. The aim was to analyse how extreme under and over taxation affects a society in order to determine whether it is possible to shrink the wealth gap without crippling the society. He implemented a flat tax system where everyone contributes 10% of their

current wealth after which the total gets redistributed equally back to all the agents. He concluded that there is a conflict between maximising total wealth and minimising inequality, however, one solution to this is to maximise the wealth of the bottom quartile. He also noted that at lower tax rates there is greater overall wealth, however, the poorest people do not share in it. When the rates are increased, the poor have a bigger share, however, the total wealth is reduced. The work of [Downey \[2012\]](#) was important for this research as it provided a framework for the incorporation of experimenting with taxation. It also gave us the idea of using our model to study progressive tax schemes which was not covered by [Downey \[2012\]](#).

[Hartshorn et al. \[2013\]](#) investigated the evolutionary emergence of particular strategies of cooperation within the framework of the Prisoners Dilemma game. This is a game in which two agents independently choose whether to defect or cooperate, and based on both of these decisions the agents would receive a reward. This forms the basis of the Prisoners Dilemma, however, they introduced an additional layer whereby agents were assigned tags which represented an identifier. They also introduced a second parameter which represented the agent's strategy. This strategy parameter would determine whether the agent chose to defect or cooperate based on its own identifier as well as the other agent's identifier. They were assigned one of four possible strategies.

No.	Name	Strategy
1	Selfish	Defects against all agents
2	Traitor	Defects against agents of like identifiers and cooperates with all others
3	Ethnocentric	Only cooperates with agents of like identifiers
4	Humanitarian	Cooperates with all agents

They applied the same methodology (with their additional strategy parameter) as that of [Hammond and Axelrod \[2006\]](#) whereby they ran 2000 cycles on an initially empty lattice grid of size 50x50. Each cycle had four phases namely: immigration, interaction, reproduction and death. The immigration phase worked as an initialisation phase where a randomly generated agent was placed at a random location in the world. It was during the interaction phase where agents would play out the Prisoners Dilemma with the other neighbouring agents and record the payoffs. From the results obtained they were able to conclude that an ethnocentric strategy emerges as the most dominant, however, there is an initial early humanitarian dominance. Further study led them to the conclusion that this initial humanitarian dominance was due to stochastic nature of the environment. This was due to the fact that initially, agents are only interacting with others of the same tag. In this situation, ethnocentric and humanitarian strategies can be considered equivalent. They also concluded that the reason the ethnocentric strategy performs better is that on the border between humanitarian and ethnocentric agents. The ethnocentric agents would defect and the humanitarian agents would cooperate. This results in a loss for the humanitarian agents and thus ethnocentric agents were more likely to survive on these borders eventually pushing the humanitarian agents farther and farther back. This approach of cooperation was useful as we were able to incorporate it into our trading model, where we were able to study the success or failure of agents using similar strategies when trading.

In conclusion, the field of Artificial Societies is relatively young and still has many applications that have not yet been fully researched. From the studies that have been done, however, we can see that there is great potential for its application in social sciences.

Chapter 4

Research Methodology

The previous sections outlined the concepts required to implement and test our model. This section contains a description of the research methodology along with the specific aspects to be modelled.

Our goal is to observe and analyse the emergent behaviour of our Artificial Society which results from the alteration of the following aspects.

Agents

- Vision
- Metabolism
- Cultural beliefs
- Number of maximum children

Environment

- Topology
- Number of resources
- Environmental upliftment

Rules

- Movement
- Information spread
- Trade
- Tax
- Investment

In order to quantify the effect of changing these various aspects we observe the society as a whole in terms of the following criteria.

- Population size
- Population distribution
- Resources available

- Average age
- Average wealth
- Wealth distribution
- Number of cultural differences
- Cultural distribution
- Number of births and deaths
- Average trade prices

4.1 Methodology

The above-mentioned aspects require a multi-agent system in order to be tested. This section outlines how this system was designed and implemented. It describes how the analysis of said system was done in order to deduce whether or not emergent behaviour had arisen and why.

4.1.1 Implementation

In the implementation phase, a multi-agent system was designed. This consisted of an episodic environment implemented as a discrete 2D grid where multiple agents were able to interact with each other and the environment. This system was implemented using Object Orientated techniques where both the agents and the environment were defined as objects. These agents were made up of different characteristics, the values of which are dependant on the particular experiment being tested. These characteristics are defined in more depth later. The agents were randomly placed within the environment at the beginning of the simulation and from then on operated autonomously. This is achieved by randomly generating numbers to represent an agent's initial x and y position. It is important to note that whenever a value is randomly defined, it has been done so using a uniform random number generator. During each episode each agent was able to interact with the environment and other agents by performing actions. The agent had a number of actions to perform per episode. Each of these actions were focused around a particular aspect; a few example actions were to collect food, reproduce, trade and spread influence. Throughout the simulation the environment was also changing based on a set of rules. One such rule was how the resources on the grid grew back after being consumed. Finally after the simulation finished we were able to analyse the information of the remaining agents along with the information generated through the course of the simulation. This information is what was used to determine the emergent behaviour and analyse the social structures that had formed. We also developed a visualiser which was used to watch the evolution of the society in real time.

To implement the multi-agent system mentioned above, objects for both the agents and environment were required. The design for both of these objects is outlined below. A controller was required which handled the operation of the simulation. This controller was used to initialise the environment (i.e. the size of the grid and the initial resource yields for each grid position) as well as the initial group of agents and where they were located on the environment. On completion of this initialisation process, the system operates autonomously. The controller was also required to be able to handle the visualiser and ensure that the visualiser mirrored the changes of the multi-agent system over time. The general flow of the simulation was that during each episode all agents were given a chance to perform their actions of moving, consuming, reproducing, trading and spreading influence. After this, the environment

performed its own tasks of regrowing the world. Once both the environment and all the agents had their turn, the episode was to be considered over and the process was then repeated for a set number of iterations.

Agents: Each agent was initialised with the parameters listed below, each of which will either be set or randomly assigned a value that is within the particular range of said parameter. For example, gender would be randomly assigned a discrete value from the set $\{0, 1\}$. It is important to note that the parameters listed below may also be predefined to a specific value rather than a random generation if particular testing requires it. For example, if we wanted to access the wealth distribution of people all with the same metabolism but differing vision values, we could set all the initial agents metabolism to a particular value and grow the society from there.

- vision: Represents the number of grid positions in each cardinal direction an agent can see per episode. It is important to note that this changes depending on the movement rule an agent is using.
- gender: Represents the sex of the agent.
- generation: This is a value which describes roughly the period of time an agent was born. The initial agents would all be considered generation 0. Children of these agents would be generation 1 and so on. The reason this is not an accurate measure of when an agent existed is because it is dependant on when its parents gave birth to it not on the number of episodes that have passed.
- resource-quantity: Represents the agent's current resource quantity or value.
- x,y: Represents the spatial location on the 2D grid.
- culture: Represents a parameter defined as an array which represents an agent's beliefs.
- age: Represents how old an agent currently is.
- history: Is a parameter defined as a linked list which stores the previous locations of the agent.
- metabolism: Represents the amount of a particular resource consumed per episode.
- initial endowment: Is half the initial amount of a resource an agent is born with. This is also the cost an agent will have to pay to reproduce.

These parameters are defined in greater detail in Chapter 5.

At the beginning of each agent's turn, they consume an amount of their resource-quantity based on their metabolism parameter. Now if the current value of resource-quantity falls below 0 the agent is removed from the simulation and in essence this agent is now considered dead. If not, the agent continues and is able to perform multiple actions during each episode of the simulation. Each of these actions can be considered a goal for the agent. These actions will be described in detail later on. Finally once an agent had survived for 80 episodes it would then have a 33% chance of being removed as if by dying from old age.

Environment: The environment works in a similar fashion to the agents, as it also contains both rules and parameters that can be randomised or predefined.

- number of agents : This parameter represents the number of initial agents in the simulation.
- number of resources : This parameter represents the number of different resources.
- size : This parameter represents the size of the grid, for all experiments we used a grid size of 50x50.

- resource grow-back rate : This parameter represents how many much of a particular resource grows back each episode.
- The following parameters are defined for each cell of the grid
 - time since last harvested : This represents the number of episodes since the last agent visited.
 - cellid : This parameter is an identifier for each cell of the 2D grid.
 - amount : This parameter represents the current amount of each resource is present on this tile.

The grow-back rule will be executed at the end of each iteration of the simulation. This is the period after all agents have first been processed and before the next iteration begins. This rule will handle the regrowth of resources across the 2D grid and will take into consideration the maximum yield of each square as well as the number of iterations since the square was last visited. This second parameter will be to ensure a more natural regrowth across the grid as it would be unnatural for a tile to regrow a portion immediately after it had been consumed.

4.1.2 Testing

The experiments were run using different input parameters for both the world and the agents. These changes include:

- Changing the number of agents.
- Increasing/decreasing the amount of a resource available.
- Changing the values of unique parameters of the agents.
- Changing the fundamental way the agents interact with other agents.
- Changing the environment the agents exist in.
- Changing the different rules which governed the agents.

An experiment can be considered as a set of 50 simulations each having 1000 episodes. All 50 simulations in an experiment are set up with the same initial structure and the results of all of these simulations were averaged in order to conduct an accurate and fair analysis. Since agents are defined with parameters, defined on ranges, each simulation is initially different. Simulations also differ because of the random initial placement of agents. However, if we looked at the population, in general, they are all initially the same across simulations. By this, we mean that the number of agents with each value for each parameter is relatively the same from simulation to simulation. This is a result of using a random number generator which follows a normal distribution. The reason we conducted 50 simulations is to limit the impact of statistical anomalies that can arise due to the inherently probabilistic nature of the simulation. The objective of running different experiments is to use different initial conditions as well as to compare different rules.

In order for any accurate analysis to be conducted each aspect was first tested independently from the others. This process was required to accurately draw conclusions from the data we generated. From this, we generated a baseline of understanding, after which we then ran simulations with aspects running in conjunction to determine a set of cross-relations between multiple aspects.

4.1.3 Analysis

The analysis was carried out by comparing the state of the simulation as well as by using the visualiser to get a better sense of how the society changed over time. This visualiser gave a real-time representation of the grid allowing for a visual medium for viewing the changing resource levels, the movement of agents as well as the evolution of the agents. For example, we could colour all the agents based on their metabolism, then run a simulation and see in real-time how the composition of agents with differing metabolisms change. This allowed us to better answer questions like do agents with higher metabolisms die out quickly and if so why. Both of these forms of analysis were done in order to determine any emergent behaviour that had developed. The use of these two forms of analysis was also a good way of cross-checking the validity of the one form with the other. A case arose where the raw data suggests one thing but was contradicted by the visualisation. These errors were the result of some implementation issues and would therefore not have been spotted if it weren't for this cross-checking approach.

The raw data is made up of aspects monitored over the duration of the simulation. These aspects are listed below.

- The change in carrying capacity over time.
- The average age.
- The average wealth.
- The amount of each resource available.
- The number of births and deaths.
- The change in the number of agents binned per gene trait.
- The change in average wealth binned per gene trait.
- The number of cultural differences.
- The change in trade prices over time.

At the end of the simulation, the data was stored in a text format. This data was then used to generate graphs which allowed for a better understanding of the Artificial Society that had evolved. This analysis of data needed to be conducted in order to provide insight into the emergent behaviour and social structures that had arisen from the evolution of the society throughout the course of the simulation.

4.2 Visualiser

For this research two visualisers were developed. One was dubbed the built-in visualiser which would operate alongside the experiments and display the movement of the agents and the changes in the environment. It was also used to visualise how the culture parameter changed during the course of a simulation. The second visualiser was a standalone system that could be utilised after an experiment had been conducted. The need for this arose due to the fact that the built-in visualiser only ran a single time and the visualisation of a particular experiment could not be rerun. For example, if you wanted to visualise the metabolisms of agents throughout the course of the experiment, you would have to change the implementation to do so and then run an experiment. It would then be impossible to also see the vision of those agents for that very experiment. It was based on this need that the standalone visualiser was developed.

4.2.1 Built-in visualiser

The figures below show the built-in visualiser at two points in time, displaying the culture parameter of an agent as a colour as well as the sugar yields in yellow. These concepts will be discussed in greater detail in the Simple Model section as well as the Culture section. The grid was generated by creating an array of tiles arranged in a 50x50 pattern to match that of the simulations. The border and background of these tiles could then be changed based on what that cell should be representing. For example, if no agent was present at a particular grid coordinate, the tile's background would be coloured, based on the sugar yield at that tile. If an agent was present, the culture parameter would dictate the colour the cell should be.




Figure 4.1: Built-in visualiser




Figure 4.2: Built-in visualiser

4.2.2 Standalone visualiser

The figures below depict the Standalone visualiser as well as how it was used to validate results.

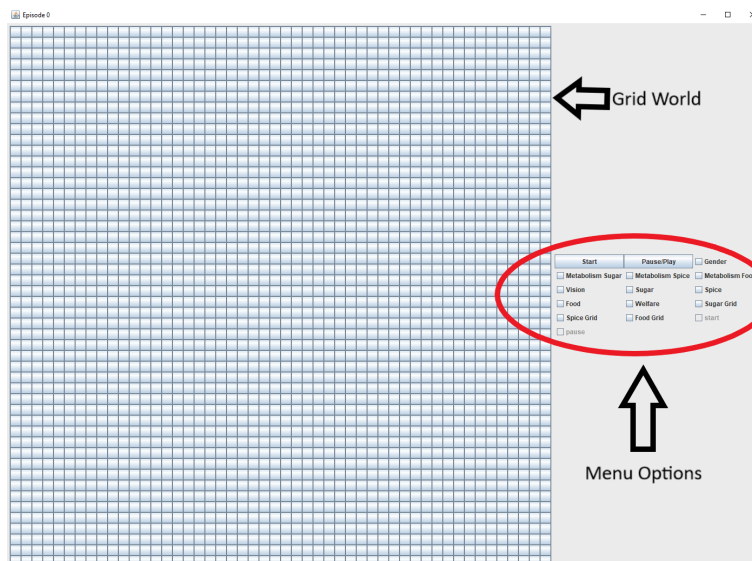


Figure 4.3: Standalone visualiser menu

Figure 4.3 depicts the main interface of the visualiser. The grid of tiles on the left was designed as previously described in the built-in visualiser. A set of check buttons were also added which can be seen on the right. These check buttons are what are used by the user to visualise different aspects of the experiment. If the user wanted to display the metabolism of the agents in the world they would simply have to press the corresponding check button. In order for this Standalone visualiser to work all the relevant information pertaining to the simulation had to be stored and then accessed by this visualiser.

This was achieved by writing simulation states to text files for each episode of the simulation. These text files would then be read in order by the visualiser and this information would then be used to display the state of the simulation. Since the state information was now stored on disk we could rerun previous simulations over again whenever we wanted to look at a different aspect.

The following figures demonstrate the visualiser in use at different times during a single simulation. It is important to note that the colouring of the tiles was just added as visual aid and is not representative of a particular value as the colouring scales with the change in values. This means that the highest values will always appear darker and the lower values will appear lighter. ¹




Figure 4.4: visualisation of sugar grid without agents

¹A working version of the stand-alone visualiser can be found at https://github.com/BrandenIngram/717431_Masters_Dissertation_ArtificialSocieties




Figure 4.5: Visualisation of agent's sugar values




Figure 4.6: First visualisation of agent's spice metabolism




Figure 4.7: Visualisation of agent's gender, blue = male, red = female

Chapter 5

The Simple Model

5.1 Introduction

Much of our inspiration for research into Artificial Societies is due to the ground-work done by [Epstein and Axtell \[1996\]](#). The Simple model serves as the basis of our research and represents our first dive into Artificial Societies. Our model is implemented as a multi-agent system in a similar fashion to that described in [Epstein and Axtell \[1996\]](#). They were able to observe the emergence of a number of interesting behaviours. The intent of our simple model is to firstly ensure we can replicate similar behaviours to those which emerged in the SUGARSCAPE. The main goal of this chapter is to simulate more complex behaviours by taking all we have learned and applying it in order to answer the following questions:

- What comparisons can be made between our model and the Predator Prey model?
- Can we generate a society with behaviour similar to that of evolution where stronger genes persist and weaker ones die out?
- Which economic trends emerge through the simple interaction between agents and the environment such as supply and demand?
- Does the way in which we traverse an environment affect our performance?
- What policies can be tested using our model which help curb contentious topics such as overcrowding or decreasing natural resources?

This model acts under the movement rule, reproduction rule as well as the regrowth rule. The algorithms for each of these are described below. The agents of this model have a singular goal and that is to seek out and collect a single resource, namely sugar, in order to survive. If they are wealthy enough they can reproduce with any agent who meets certain criteria. The cycle of the simulation is as follows:

For each agent do:

- Seek out the tile with richest resource quantity within vision of the current agent.
- Move current agent to that tile.
- Remove all resources at that tile and give them to the current agent.
- Consume metabolic requirement of resource for current agent.
- Reproduce if current agent meets requirements.

- Check to see if the current agent is dead.

This process repeats for 1000 episodes and the relevant information is recorded for future analysis. Each episode is considered to be a year in terms of the agents. Finally once an agent turns 80 years old, or in other words, has existed in the world for 80 episodes, that agent has 33% chance of dying each episode as if by old age.

As previously discussed this multi-agent system was made up of agents with internal parameters and rules. The key internal parameters and rules that were used in the simple model are outlined below.

Internal Parameters

- Metabolism - Determines how much sugar is consumed per episode. Initialised randomly as a value in the set {1, 2, 3, 4}.
- Vision - Determines how much of the world an agent can see (the number of tiles in each direction). Initialised randomly as a value in the set {1, 2, 3, 4, 5, 6}.
- Gender - Whether male or female. Initialised randomly as a value in the set {0, 1}.
- Sugar - How much sugar an agent currently has. Initialised randomly as a value in the range [5, 25]. The reason for this initial endowment of sugar is to ensure our population does not die instantly.

Rules Movement Rule 1 works as follows:

- For each agent do:
 - Look out as far as their vision permits using a Breadth First Search.
 - Determine which vacant tile has the highest yield. If multiple options are available choose to go to the closest one.
 - If the highest yield is zero, randomly pick a tile to go to.
 - Determine the path to that tile.
 - Return the adjacent tile to the agent which is along that path.

For this rule, you would see agents moving from one tile to the next adjacent tile on their way to their goal. It is also important to note that this search occurs on each episode, so there is a very high likelihood that the goal will change from episode to episode. This was to ensure agents don't chase after rewards that have already been collected by other agents.

Movement Rule 2 works as follows:

- For each agent do:
 - Look out as far as their vision permits using a Breadth First Search.
 - Determine which vacant tile has the highest yield. If multiple options are available choose to go to the closest one.
 - Return that tile with the highest yield.

In this second movement rule, an agent will move directly to the location with the highest yield.

Movement Rule 3 works as follows:

- For each agent do:

- Initialise a “maxsugar” variable to be a minimum.
- For each vacant tile do:
 - * If the sugar value at the current tile is larger than our “maxsugar” variable. Change “maxsugar” to be the current sugar value and record this location for future.
- After each tile has been processed return the location of the “maxsugar” variable.

Using this movement rule the agents would teleport like in the previous rule directly to their destination, however, in this case, they would be able to teleport anywhere in the world.

The agents do not choose which movement rule they operate under but will rather be decided upon beforehand depending on the particular experiment in question.

The Reproduction Rule works as follows:

- For each agent do:
 - Look out as far as their vision permits using a Breadth First Search.
 - Check to see if an agent is seen.
 - If another agent is seen check to see if they meet all criteria for reproduction.
 - If all criteria have been met, check to see if there is space adjacent to either the current agent or the potential partner for the new agent to be born.
 - If so create a new agent at one of the vacant adjacent tiles.
 - Reduce the current agent’s and the partner agent’s sugar value based on the initial endowment the child requires.
 - Give this endowment to the child.

The criteria required for reproduction are :

- Agents need to be of opposite sex.
- Both agents must be older than 16.
- Both agents must be younger than 50.
- Both agents need to have enough resources to contribute to the initial endowment.

The initial endowment of the agent is the initial amount of sugar they are born with in order to survive. Each parent contributes half of this amount and is based on what they originally received. For example, if ParentA had an initial endowment of 50 they would contribute 25 to their child. This initial endowment can be considered the cost of raising a child on the parents. The new agents would also have the same internal parameters of their parents. The values for each of these were derived from their parents with a potential occurrence of a deviation away from either parent. For example, the vision of a newborn agent had a 45% chance of coming from their father, a 45% chance of coming from their mother and a 10% of being defined as a value in {1, 2, 3, 4, 5, 6}. This does still mean the value could be the same as either parent but, it also allows for a different value to occur.

The environment also has internal parameters and a rule, which is outlined below.

Internal Parameters

- Size - Determines the size of the grid.
- Regrowth-rate - Represents the number of episodes a tile would take, after being harvested, before starting to regain sugar yields.
- Regrowth-amount - Represents the amount of sugar a tile would replenish after the sufficient delay.
- Capacity - An array which kept track of each tile's highest possible yield.
- Current - An array which kept track of each tile's current yield.
- Last-visited - An array which kept track of when each tile had been previously harvested.

The values and ranges of all of the parameters listed above have been chosen in line with the previous research done by [Epstein and Axtell \[1996\]](#) as well as from considering global census data from [Factbook \[2015\]](#) where for example we can determine that life expectancy for developed countries is roughly 80 which is in accordance with our experiments. In [Figure 5.1](#) below we have a depiction of the sugar levels across the world, the tiles with darker shades of yellow representing higher yields and the white tiles having no yield. The darkest yellow has a yield of 4 going down to white having a yield of 0. This will be the same arrangement of resources utilised for all experiments to follow.




Figure 5.1: Visualisation of the replica of the SUGARSCAPE model depicting the levels of sugar

Rules

How the Regrowth Rule works is as follows:

- For each tile do:
 - Check to see when the tile was last visited by an agent.
 - If the tile was last visited longer than the regrowth delay then do:
 - * Increase the sugar quantity by the regrowth rate.
 - * If the value exceeds the capacity, set the sugar quantity to equal the capacity.
- Increase the time since last visited.

Through the implementation of these rules, we are able to generate a fully autonomous society, where agents interacted with both the environment and each other. These interactions result in the emergence of various phenomena. These being population dynamics, evolution as well as phenomena found in Economics, all of which will be discussed in the sections to follow. ¹

5.2 Predator Prey Model Comparison

When experimenting with and comparing our model to the SUGARSCAPE, we observed an interesting dynamic behaviour between the resources available for consumption and the population of agents. That is when more agents existed, fewer resources were available. The inverse was also true for when there were fewer agents. It was because of this observation, that we decided to compare it to the Predator Prey Model which represents a similar relationship. Therefore in this section, we compare the similarities and differences between the two models.

In order to conduct a comparison, we must first understand what the Predator Prey Model is and how it works. Suppose there exist two species within an ecosystem, where the first species only eats the vegetation and the second feeds off the first. Here the second species is the predator and the first is the prey. For example, foxes hunt grass eating hares, here the fox is the predator and the hare the prey. The Predator Prey Model is a system of equations which represents the relationship between these two actors.

We shall generate our own Predator Prey Model as in Zill [2016]. Let $x(t)$ and $y(t)$ denote the populations of the foxes and rabbits at time t respectively. They state that if there were no rabbits, one might expect that the foxes, lacking an adequate food supply, would decline in number according to the equation below.

$$\frac{dx}{dt} = -\alpha x, \alpha > 0 \quad (5.1)$$

Here α represents the death rate of foxes and is proportional to the number of foxes present. When rabbits are present in the environment, it is reasonable that the number of interactions between the two species would be jointly dependant on both populations, that is proportional to the product xy . By

¹A full visualisation of this simple model can be found at https://github.com/BrandenIngram/717431_Masters_Dissertation_ArtificialSocieties

adding rabbits we introduce a supply of food for the foxes, so foxes are added to the system at a rate of βxy where $\beta > 0$. Adding this rate to our previous equation gives us the following.

$$\frac{dx}{dt} = -\alpha x + \beta xy \quad (5.2)$$

Next, we consider the population of rabbits. If there were no foxes and under the assumption of unlimited food supply to forage on, rabbits would grow at a rate proportional to the number of rabbits at time t

$$\frac{dy}{dt} = \delta y, \delta > 0 \quad (5.3)$$

Again by introducing foxes a model for the rabbit population would decrease at a rate at which rabbits are eaten by foxes. This can be defined as σxy where $\sigma > 0$ giving us the equation:

$$\frac{dy}{dt} = \delta y - \sigma xy \quad (5.4)$$

This gives us a system of equations which defines our Predator Prey model.

$$\begin{aligned} \frac{dx}{dt} &= \beta xy - \alpha x && \text{(foxes)} \\ \frac{dy}{dt} &= \delta y - \sigma xy && \text{(rabbits)} \end{aligned} \quad (5.5)$$

The figure below was generated by solving the system above with the following initial conditions:

$$\begin{aligned} \alpha &= 0.2 & \beta &= 0.07 \\ \sigma &= 0.8 & \delta &= 4.7 \\ x(0) &= 4 & y(0) &= 4 \end{aligned} \quad (5.6)$$




Figure 5.2: Predator Prey model

From Figure 5.2 above we can draw conclusions of how the populations react with respect to one another. For example, after a period of increase in the prey population due to reproduction, we see an increase in the predator population as the predators now have a source of food required for their survival. However, this increase in the predator population leads to a decrease in the prey population as they are now being hunted and killed. During this period when the prey population is decreasing the predator population is still increasing as there is still sufficient food to support the growth. Once the prey population reaches sufficiently low population numbers to no longer support the predator population, we see that the predators begin to die out from starvation. This decrease in predators allows for the prey population to increase from normal reproduction. This cycle of dependent fluctuation continues for both populations until $t = 20$

In order to draw comparisons, we should try to set up both models in similar fashions in order to see if a correlation exists between the two. Therefore, the initial population of Foxes will be the same as that of the agents, 200, and the initial population of rabbits will be the initial number of tiles which contain a sugar value greater than 0, i.e., a tile which is available for consumption. This system of equations is defined below.

$$\begin{aligned}
 \alpha &= 0.377 & \beta &= 0.0002 \\
 \sigma &= 0.002 & \delta &= 0.42 \\
 x(0) &= 200 & y(0) &= 2000
 \end{aligned}
 \tag{5.7}$$

Solving for this system of equations results in the figure below.




Figure 5.3: Predator Prey model realistic conditions

From Figure 5.3 above we see the same relationship that was discussed above. The right-hand side of Figure 5.3 is when we plotted the predator population versus the prey population.

We also note that through interactions in our multi-agent system we were able to see that the emergent behaviour that evolved was similar to the Predator Prey Model. In our case, the agents can be considered the predator and the sugar in the environment, the prey. We observe that when the population of agents increases, the total level of sugar decreases. This is because there is more consumption of sugar occurring at higher population sizes. If we observe a population decrease, this implies that there are fewer agents consuming sugar and therefore the sugar has greater opportunity to regrow to higher levels. This

relationship perfectly matches that described in the Predator Prey model. The figures below backup this observation with empirical evidence.




Figure 5.4: Predators (population of agents) and Prey (amount of sugar) comparison

In Figure 5.4 above we can easily see the correlation between the population and the total sugar available, or rather our prey versus predator populations. For example, as the population is decreasing in the first 20 episodes the amount of sugar is increasing. The initial sudden drop off in sugar is due to the introduction of the initial 200 agents. We can then note that as the population starts to grow from this point so, in turn, does the amount of sugar decrease. This relationship can be seen to continue throughout the course of the simulation. Therefore, much like before we can see that when the population of agents increases this leads to a decrease in the amount of sugar available to our prey population. Finally, we also see that when the population of agents decreases this leads to an increase in the amount of sugar available.

There are, however, important differences to be observed. In the actual Predator Prey Model described previously there exists a delay between when we see the prey population decrease as a result of the increasing predator population. The same can be said when the predator population is decreasing, we see a delay in the increase of the prey population. On the SUGARSCAPE we see no such delay, as when our predator population increases, immediately our prey population decreases. When our predator population decreases immediately we see an increase in our prey population. Based upon this we concluded that our model does not perfectly correlate to the behaviours observed in the Predator Prey Model, but even though these are two vastly different models we are still able to obtain similar qualitative behaviour.

5.3 Survival of the Fittest

In previous sections, we described how agents were made up of internal parameters which affected how the agent performed within the world. For example, in the simple model, the metabolism parameter affected how much sugar an agent consumed per episode. These parameters share a special relationship with the reproduction rule, as when new agents are born they inherit the value of these parameters from either their father or their mother. We have also included an opportunity for a genetic mutation to occur whereby the value of the parameter that the child receives is from neither parent. The rate at which this

occurs is 10%. For example if ParentA has a metabolism of 1 and ParentB has a metabolism of 2, the child of AB would have a 45% chance to be born with a metabolism of 1, a 45% chance to be born with a metabolism of 2 and 10% chance of mutation. When this mutation occurs there is a uniform chance of being born with a metabolism of 1, 2, 3 or 4. The same approach has been applied to all other of the internal parameters of an agent. This approach means that the level of diversity depends on the initial population and for that reason, we start the simulation with a large population of which each agent's internal parameters have been assigned with randomly generated values across fixed ranges. It is, therefore, the goal of this section to analyse how these factors affect the structure of the overall population in terms of their internal parameters as well as how that structure changes over the course of the experiment.

The behaviour that emerged from the course of these experiments across different simulations was that of a Survival of the Fittest model or a representation of Natural Selection as described from Darwinian evolutionary theory. [Depew et al. \[1995\]](#) describes Natural Selection as a process which acts to preserve and gain small advantages from genetic mutations. For example, say a member of a species grew webbed feet and thus gained an advantage from being able to swim faster. The offspring of this member would inherit this advantage and would, in turn, pass it along to its own offspring thus perpetuating this advantageous mutation. The members who did not have this mutation of webbed feet would, therefore, be at a disadvantage and on average would perform worse in the environment this species exists in. These disadvantaged members would gradually die out over a period of time. Natural Selection is the natural equivalent to selective breeding whereby humans select partners which have characteristics that are beneficial in the hope that their offspring inherit these traits.

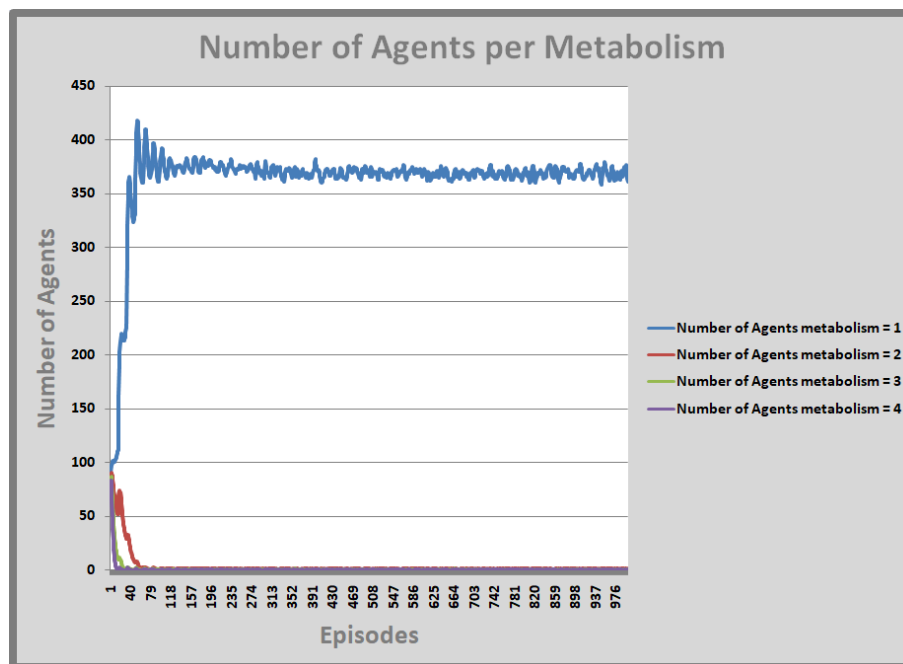


Figure 5.5: Number of agents grouped by metabolism for the simple model

Figure 5.5 was generated by keeping track of the number of agents per each metabolism over the course of the 1000 episodes and then averaged over the 50 simulations. We note that the initial population contains a relatively equal amount of agents with each metabolism which shows that there is equal opportunity for any of the agents to succeed. From this initial point, we can see that the population quickly becomes weighted towards a single metabolism. The agents with a metabolism of 4 are the

first to become extinct followed by agents with the metabolism of 3 and then finally agents with the metabolism of 2. Whereas agents with a metabolism of 1 exist in far greater numbers tending towards a value between 350 and 400. It is important to note that due to the chance of mutation a gene such as metabolism can reappear with any of the values 2, 3 and 4.

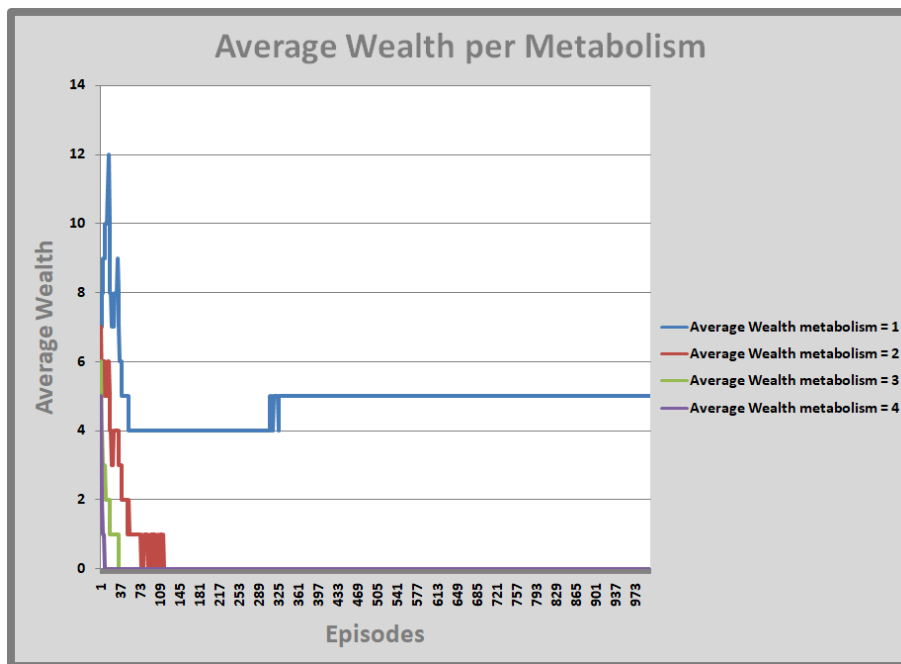


Figure 5.6: Average wealth of agents grouped by metabolism for the simple model

Population size alone cannot be used to determine if a particular trait such as a low metabolism is beneficial. However, if we examine Figure 5.6, where we consider the amount of sugar an agent has to be their wealth and then determine the average wealth per metabolism. We can conclude that agents with a metabolism of 1 are indeed the fittest or most successful due to the fact that these agents have the highest average wealth.

Next, we consider the vision attribute. Now just by considering how vision works, where if you can see further you should have a greater opportunity of seeing high yield locations of sugar, one might assume that agents with a higher vision perform better.




Figure 5.7: Number of agents grouped by vision for the simple model

If we generate a graph for vision in a similar fashion for metabolism, we get Figure 5.7 which resembles the assumption made above. Here we can see that the number of agents with a vision of 6 booms massively to around 400 after which this growth rate settles. We noted that the lower the value of the vision parameter the lower the population of agents with that vision value. This implies that the relationship seen for metabolism also emerges in vision, whereby agents with the better more efficient genes for the environment appear in far greater numbers.




Figure 5.8: Average wealth of agents grouped by vision for the simple model

Once again when we look at the average wealth per vision in Figure 5.8 above, a similar picture can be seen. This picture is that agents with higher vision have a greater average wealth. Notably, agents with vision 5 periodically have an average wealth greater than that of agents with vision 6. This results from what can be seen in the population graph where an agent with vision 5 is born due to a mutation occurring. This single agent will be relatively wealthy initially from its initial endowment, resulting in the spike in the graph. Agents with a vision of 6 have a stable average wealth above 0 whereas all other periodically drop to zero. Again this is due to all agents with those parameters periodically becoming extinct. In conclusion, as a result of the reproduction rule and the carrying down of genetic traits from generation to generation, we can see that the society which emerges is one that operates similarly to that described by Natural Selection. In other words, the agents within the society evolve to be the most efficient and successful genetically for their environment. A question can be asked, at which point does vision no longer benefit an agent. This question is answered in Section 5.5 where we compare movement rules. One rule analysed is that of a global movement, which can be considered the highest possible vision.

An important difference can be noted in the implementation outlined above to that of the definition of Natural Selection. This difference is that in Natural Selection the population starts off with all members of the species having the same genetic traits, from this point a mutation occurs when a member diverges from this uniformity ultimately leading to the best traits emerging. However, in this implementation, not all members of the species start off the same. In order to simulate something more akin to the definition of Natural Selection we implemented a system whereby everyone's metabolism was initially defined as 3, from this uniform initial point there was 10% chance that the child of an agent would mutate whereby they would either get a value 1 higher or 1 lower to that of either parent.




Figure 5.9: Number of agents grouped by metabolism for the simple model with the more realistic mutation

Figure 5.9 shows that initially the population is made up of agents all with a metabolism of 3. We have previously shown how this is a relatively poor performing value for this trait. As a result of that, the population decreases very quickly in accordance with what we have seen before. Another observation is that very early on we only see agents with values either side of 3. We then start to see an increase in

the population of agents with a metabolism of 2 which is a far better value than that of 3. The increase in this population allows for a mutation to occur which results in agents with a metabolism of 1, the best performing value of the previous experiments. The population of these agents then booms out very quickly until it converges to a stable level just under 400. We can also see that the population of agents with metabolism 2 do not die out but rather survive at low levels and this is due to an agent with parents of metabolism 1 being born with a mutation resulting in a value of 2. In conclusion, here we can see a society which initially starts off uniform, and through a random chance of mutation evolves to a relatively uniform society whose traits are the most efficient and best performing for the environment. This has also demonstrated that although the population started off initially different to that of what we saw in Figure 5.5 the society evolves to the same end point where the best genes are the ones to have survived.

In conclusion, we have demonstrated that agents which have better genetic traits not only perform better but also due to this better performance have a greater chance of carrying down their strong genes. This behaviour is in accordance with that described by Darwin so many years ago whereby a species will evolve to a dominant strain of genes. In our case, this means having the lowest metabolism and greatest vision. This is important as it shows how our model can represent this complex real-world behaviour of evolution

5.4 Supply and Demand

In this section, we look to study the relationship that exists between the demand a population has for resources with the supply of resources that the environment provides. We have already shown that population sizes tend to converge to stable levels. The question that remains is what role does the environment have in determining that level. In order to determine this, we generated a graph which compared supply versus demand. By taking the Total Metabolic Demand of the society (the sum of all agent's metabolism) and comparing that to the Total Sugar Available per episode.




Figure 5.10: Total sugar available and the total metabolic demand comparison for the simple model

After analysis of Figure 5.10 above the following observations and conclusions were made. At the

beginning of the simulation, we see a sudden drop in both supply and demand, which is due to the environment now having agents existing on it. From this point, a relationship can be seen whereby as the total metabolic demand increases the total supply of sugar decreases. Conversely when the total metabolic demand decreases the total supply increases, thus forming an inversely proportional relationship.

To conclude, population sizes have a direct effect on the abundance of resources available. Taking this relationship to its extremes would mean that the environment would become void of any resources. This situation could then result in the suffering of a population and possible death of its members. Furthermore, the strain on one society could result in strain on another whereby people are forced to move to obtain resources. It is for these reasons that studying the effects of population on an environment is important. This relationship, that emerges from the interaction of the agents collecting sugar as well as the environment regrowing, led us to believe that supply and demand laws might be at play. In Microeconomics supply and demand is an economic model used to determine price in a market. A deeper look into Supply and Demand and how that affects trade prices of goods will be explored with respect to trade in Chapter 7.

5.5 A comparison of movement

In the previous section, we have analysed how a population evolves to have the best genetic traits possible in terms of their internal parameters. In this section, we delve into what behaviours would be the best genetically. More specifically, we focus on how agents traverse the environment and whether or not that has an impact on their performance. This is important as if it does have an impact, this would mean that not only are your internal genetics important but also the way you act within the environment. The way in which an agent traverses an environment can also be considered a measurement of intelligence. For example, agents who move better can be considered smarter than those with worse performing techniques. As the Movement Rule handles how an agent traverses an environment we decided to implement two additional variants to this rule. This was done in order to achieve the goal of analysing the performance of different approaches to movement.

In Movement Rule 1 agents look out in their vision to locate the highest yielding tile. They then move one adjacent block at a time to that location, however, each turn they reevaluated which tile to travel to. This is in case this tile is no longer the best tile to travel to. Along this journey agents still collect resources at each tile they come to. Movement Rule 2 is the same rule used in the previous experiments and entails locating the highest yield tile within the agent's vision and instantly teleporting to that location. Movement Rule 3 is implemented such that an agent locates the highest yielding tile on the entire grid and then teleports directly to it. This can be considered to be Movement Rule 2 with a global vision. Three experiments were conducted each only allowing its population to perform under a single rule, the results are shown in Figure 5.11 below.




Figure 5.11: Total number of agents for each movement rule

In Figure 5.11 we observe that using Movement Rule 3, or the global move, resulted in the highest population, roughly double that of either of the other two. In all three we noted the cyclic rise and fall of population sizes, however, with rule 1 we saw a flattening of these cycles. We also recognised that for rule 3 the periods where the population is decreasing are very sudden and harsh, followed by an equally sudden boom. We concluded that this was due to a sudden dying off of a population which had grown older together. The reason they grow old together is as result of only having a small chance of dying from not being able to find food since they can teleport. The agents who remained were also wealthy from an albeit not as long a life, but one where they were able to gather surplus wealth. Now since the population died off there was also a lot more vacant tiles where new agents could be born. These gaps are then suddenly filled by the reproduction between these rich agents and then the cycle repeats again when they, in turn, are the old ones. The same logic is true for rule 2, however, it is less extreme.




Figure 5.12: Average age of agents for each movement rule

In Figure 5.12 we again see that with rule 1 the cycles flatten out. By both these graphs flattening out we can conclude that the birth and death rates are equal, moreover that births and deaths are not

happening in large clusters. We also noted that there are large oscillations in the average for both rule 2 and 3. These oscillations are, as previously discussed, due to the clustering of deaths and births in short periods. These clusters lead to a population that ages at the same time. Therefore from a large number of births in a short period we see that later on, a large number of deaths occur in a short period. The danger that arises is when we see the average age become too large, and not enough new agents are being born to support the next generation. This has the potential to result in extinction level events for the population.




Figure 5.13: Average wealth of agents for each movement rule

In Figure 5.13 the most interesting observation made was of how low the average wealth for rule 3 was. This observation becomes clear when considering the fact that if we increase the population of an environment they still compete for the same amount of resources as that of a smaller one. What we actually see here is that the population maintains itself at such a large level, that almost all of the agents living in that world are just barely scraping by with an average of 1 sugar (wealth). We can also conclude that rule 3 is the most efficient rule for consumption of an environments resources as it manages to maintain a level of consuming everything there is to offer for the duration of the experiment. We can also see that both rule 1 and 2 produce equal results for their respective wealth. In conclusion, we deduced that since rule 1 results in slower movement across the world. This, in turn, allows the environment a chance to regrow and is also the reason we do not see a spike in the wealth graph.

The lack of spikes in the wealth graph also shows why we do not see a population explosion since there aren't a large number of agents with excess sugar which will be used to reproduce. What emerges from agents using movement 3 is that all the agents are now competing with each other globally rather than local regions determined by their vision. This global competition makes it harder for individual agents to consistently outperform others. For example in rule 1 or 2, if an agent is located in a region with high yields of sugar they only have to compete with a small portion of those around them and those who will eventually reach this region as well. This allows a greater chance of generating large amounts of wealth. In the case of rule 3, everyone in the entire world is effectively within that small portion. An argument can be made that the agent now has more tiles to go to and therefore it wouldn't matter, that they are sharing them all. However, since this rule results in an increased population, you are sharing all those tiles with far more agents than if the world was acting under rule 1 or 2. In this experiment, we observed how the society that evolves for rule 3 is one which has perfectly optimised for maximising the number of agents alive. Now, this doesn't mean that it is necessarily the best policy for a real-world situation.

This is due to the fact that quality of life becomes an important issue, and we have shown how under rule 3 the quality of life is very poor. Figure 5.14 below was generated to validate our claims where rule 1 resulted in the agents consuming the resources of the world slower. The reason this validates our claim is that we see a consistently higher level of sugar available under rule 1. One could argue that this is a result of population size differences. If we compare rule 1 and 2 in Figure 5.11 the population size of rule 2 is consistently lower beyond episode 400. However, in Figure 5.14 the supply for rule 2 only exceeds that of rule 1 for a very short period of time.




Figure 5.14: Total amount of sugar for each movement rule

In conclusion, we have shown that by changing the way agents operate within an environment we can see that the society that emerges can have vastly different behaviours. These behaviours being how efficiently they consume the resources of the world as well as how this resulted in different styles of reproduction. Whether it be the large batch reproduction and death of the very efficient rule 3 or the more drawn out reproduction resulting in no batch deaths. It is important to note that for all the previous experiments in this section, the entire population in each world had the same movement rule. If we were to do an experiment where we initialise an agent's movement rule to be one of the three described, in the same fashion as vision, we could expect to see that the population which emerges to all be working under Movement Rule 3. This is because agents using this rule would be vastly outperforming all others in most cases and in the worst case would be performing at the same level.

5.6 Birth Control

The one-child policy was a population planning policy implemented in China which formed part of the family planning policy [Goh 2011]. It was introduced in 1979 and began to be formally phased out in 2015. It is important to note that the fertility rate had already begun to decline around 1963 prior to its introduction. The Chinese government, however, observed the global debate over a possible overpopulation catastrophe suggested by organisations such as Club of Rome and Sierra Club. These organisations deal with a variety of international issues, including the world economic system, climate change, and environmental degradation. Song Jian, a top Chinese official read two influential books on the movement, *The Limits to Growth* [Meadows *et al.* 1972] and *A Blueprint for Survival* [Goldsmith 1974] while visiting Europe in 1979 [Hvistendahl 2010]. Inspired by these works and together with a

group of mathematicians, Song determined the correct population of China to be 700 million. A plan was therefore prepared in order to reduce China’s population to the desired level by 2080. The one-child policy would be one of the main instruments of this social engineering. In spite of some criticism inside the party, the plan was officially adopted in 1979. The plan called for families to have only one child each. The goals of this policy were to curb a then-surging population and limit the demands for water and other resources, as well as to alleviate social, economic and environmental problems in China. On September 18, 1980, the policy was formally implemented as a temporary measure in the hope it would achieve the goals outlined above. This section will demonstrate the effect of enforcing such a policy and others like it on a population.

In order to simulate a similar policy in our world, an additional parameter was introduced to the agent in order to keep track of the number of children that agent had produced. We then ran the same simulation setup as before whereby we ran 50 simulations with 1000 episodes each, with the simulations acting under the movement, reproduction and regrowth rules. However, one slight change was made to the reproduction rule, whereby we ensured that a parent would not exceed a predefined number of maximum children. We ran this for a range of different values for the maximum number of children. The results can be seen below.




Figure 5.15: Total number of agents per the differing birth control schemes

From Figure 5.15 above when the agents are limited to a maximum of one child, the overall population decreases to a point where extinction occurs. This result was expected as by limiting two parents to one child each you are at best halving the population periodically until a point where everyone has died. Secondly when the max child is set to 3 or 4 as well as when no restriction is enforced we see that the population initially explodes up to a certain level after which it converges to a stable population. This stable level is highest when the maximum number of children was set to 3 and 4. This shows that an introduction of a maximum child policy can have a very prominent effect on how a population changes over time. Finally, when the max child was set to two we saw a notably slower population boom in the earlier stages indicating that the population was growing slower due to this reproduction restriction. The population does even converge to a higher level than that when no restriction was enforced, it just takes

a longer period of time. The last observation to make is that the “no control” policy converges to a lower level than that of the other 3 successful policies. This is due to the fact that by having no restriction agents reproduce whenever they can and therefore incurring the cost of reproduction more often than that of other strategies. What we find is that agents reproduce too much when they are wealthy and end resulting in their death as they can no longer find enough resources within their vicinity to support themselves.

Just analysis of the population size is insufficient to determine which policy benefits the world the greatest. For this, an in-depth analysis of the average wealth is required.




Figure 5.16: Average wealth of agents per the differing birth control schemes

Firstly from Figure 5.16 above, when the max child is set to 1 the average wealth grows larger and larger in conjunction with the decreasing population, indicating that there are a few really rich agents alive. The average wealth then plummets. This is due to the fact that the population is made up of very few rich agents all of which can no longer reproduce, they then eventually all die out from old age. The most important observation here to note that when the max child was set to 2 we see that the average wealth converges to a higher stable level than that of any other policy. This level is roughly double that of the rest.

In conclusion, we can see that when the maximum number children were set to 2, not only does the population size reach relatively similar levels, even outperforming the no control policy, but it, more importantly, results in a higher average wealth for its population. This means that instead of the world supporting a high population where everyone is poor and can be considered to have a worse quality of life, we could implement a policy which has a positive effect on the agents’ wealth without drastically affecting the population size.

In all the previous experiments, the population was initialised with different restrictions on the number of children an agent could have, however, this is unrealistic in a real-world situation where our populations are already evolved to some degree. For this reason, instead of introducing a policy at the beginning of the simulation, the agents are allowed to have as many children as they are capable of. This is the same

as the no control policy seen in previous experiments. At episode 500 we introduce a 2 child policy, thus limiting all currently living agents and all new agents to come to having a maximum of 2 children. The results of this experiment can be seen below.




Figure 5.17: Change in population size with introduction of 2 child policy at episode 500




Figure 5.18: Total Births and Deaths with introduction of 2 child policy at episode 500




Figure 5.19: Average age and wealth with introduction of 2 child policy at episode 500

In Figure 5.17, an immediate effect of the change in policy can be seen at episode 500. Before this event, we see the same population oscillations as seen before. After episode 500 we see a period of chaos followed by a period of population decline. This decline is no worse than the bottom periods of the population oscillation. The reason for the decline is partly due to the fact that the population was ageing thus leading to fewer births and more deaths as well as the fact that the policy limited the number of births. However, afterwards, we see a steady growth in population which converges to a stable level around 340. It is also important to note that there are no longer any oscillations in the population. This is due to the world no longer experiencing population explosions as a result of too many births in a short period of time, which can be seen in figure 5.18. The rate of births becomes linear and not periodic after the advent of the policy change. When a population explosion (a large number of births in a short period of time) occurs, this results in the majority of the population all ageing together. This means the average age will also oscillate periodically and when the population gets older we see fewer births. These oscillations are dangerous the steeper they get, as they show how a population can become too old to generate new children and then the older population dies out, there are no longer enough young agents to sustain life and regrow the population. For example, during this period of population decrease, the survivability of the population would be at higher risk of extinction due to disaster. Therefore, a policy which reduces this oscillation while maintaining sustainable wealth levels would benefit a society. In figure 5.19 the average age after a period of time converges to a stable level around 38, and our average wealth also converges to a stable point which is well above the minimum level of survival. This experiment shows that even when introducing a birth control policy into an already developed or evolved society will have an effect on population dynamics. It has also been shown that the quality of life of the agents under this new policy is better and more sustainable to that of when no control is present.

It is important to note that [Chen \[2003\]](#) states that there is evidence showing population decline was not only as a result of the Chinese policy. This includes disparities between the number of males to females as well as other pre-existing policies. These become a problem when measuring the effect of the One Child Policy. In our research, however, we were able to isolate the effect of differing population control policies by leaving all other factors constant. This section demonstrates that independent on when a policy (such as the one the Chinese implemented) which controls births is introduced, it will have a major effect on the population size, average age as well as the average wealth of its members. With

overpopulation being a debated issue of concern, we might see more of these types of policies being introduced. However, when it comes to government implementation and enforcement of these policies, there is also a question of individual freedoms that need to be addressed.

5.7 Conclusion

The goal of this chapter was to implement and test the validity of our Simple Model as well as demonstrate its ability to simulate the emergence of complex behaviour. In our pursuit of this end, we established evidence of similarities between our model and that of the well known Predator Prey model. There were also similarities seen in the economic model of the relationship between supply and demand, where our demand was based on the metabolisms and supply based on the resources available. This observation gave a sense of validity to the results obtained as well as future results to be presented in the upcoming chapters. Another interesting behaviour that emerged through the course of the experiments was that of evolution. We were able to observe the emergence of a dominant gene strain which emerged from the success or failure of agents. This behaviour matched perfectly to the behaviour described by [Depew *et al.* \[1995\]](#). Where evolution was concerned about internal parameters, our analysis of movement involved the effect of different fundamental behaviours. Here we saw how different approaches to how agents traversed the environment affected the efficiency and overall success of agents. The most successful being those agents who were able to move to the best location anywhere in the world. The final observation we were able to make was that of how drastically population control mechanisms have on a society. Here we saw how by restricting the number of children agents could have, resulted in greater population sizes, longer living agents as well as average higher wealth overall. This chapter has shown that our model does indeed serve as a strong staging point, which will be expanded upon in the following chapters to simulate more complex behaviours.

Chapter 6

Culture

6.1 Introduction

The Simple Model chapter dealt with how a population of agents evolved to best suit their environment through the emergence of particular genes, which if present in an agent would result in higher success rates. This evolution was as a result of the direct interaction between agents which occurred in the form of reproduction. This interaction had a direct effect on the agents in question because the process of reproduction came at a cost, which represented the cost of raising a child. This chapter revolves around this idea of direct interactions having a direct effect on the interacting agents. More specifically it introduces a new rule (The Influence Rule) and internal parameter (Culture Parameter), as well as how they can be used to model the effect of spreading ideas through one-on-one communication between people or agents. [Epstein and Axtell \[1996\]](#) originally introduced the idea of using a cultural parameter to represent an agent's internal cultural factors, which could then be used to change their behaviour in relation to other rules. This parameter was implemented as a sequence containing only zeroes and ones. Each agent's cultural parameter was randomly initialised at birth. This cultural parameter was of equal length for each agent and each element of the set was called a tag. An example of the cultural parameter would be 100010101101. [Epstein and Axtell \[1996\]](#) also introduced the idea of a cultural transmission rule which defined how agents could change each other's cultural parameter thus leading to the change in the distribution of tags in the society over time. The cultural transmission worked as follows:

Suppose there are two agents A and B and that A is a neighbour of B. An agent is a neighbour if they are on adjacent tiles to one another in the four primary directions namely up, down, left and right. A tag position is randomly selected, then both agents compare their value at that position. If A has a value of 1 and B has a value of 0, this is considered to be a disagreement. Agent B then has his tag flipped to match that of agent A. If they both have the same value then it is considered that both agents agree and no change or flip is made. This process is then repeated for all neighbours of A, after which we consider the turn of A to be complete. We can consider this interaction between the two agents to represent a dialogue or discussion on their beliefs. This entire process then repeats for all agents in the world. Since this same process is repeated for all agents the agent B in the above explanation will can be considered as agent A when their turn comes around. This means that both agents will get a turn to influence each other over the course of an episode. Cultural groups were then defined based on the number of zero tags

in the agent's cultural parameter. ¹

Agent Group	Number of Zeros
Blue	0-3
Green	4-7
Red	8-11

Table 6.1: Cultural Group Breakdown

Based on Table 6.1 above an agent with a cultural parameter of 111001101100 would be grouped as Green. By visualising the cultural groups of their agents over the course of a simulation, Epstein and Axtell [1996] observed that the society converged to a uniform cultural group through the one-on-one interaction between two agents.

In order to simulate the spread of knowledge or culture in our own model, we decided to implement a similar system to that of Epstein and Axtell [1996]. However, we extended the range of our cultural parameter by defining it as an array of size 9 where each element in the array could have an integer value on the range [0, 3]. This implementation was adopted as we wanted to see whether or not the same results would hold when we expanded the size of the influence parameter from that used by Epstein and Axtell [1996]. Our Influence Rule was also implemented differently whereby in our case an agent's neighbours are all other agents within their own vision. This means that agents are interacting with not only the agents in the four primary directions but also those within the circle of vision around them. This was incorporated to make the Influence Rule operate in a similar fashion to that of the Movement Rule. We also incorporated a chance in which no change in tags would occur. Thus the modified rule worked as follows. If A sees B in their vision he/she enters a "dialogue" with B. A then compares each element in his/her culture array with that of B if there is a difference, there is a 10% chance of A changing his/her value to be that of B and 10% of B changing his value to be that of A. This process happens for each agent within A's vision and occurs every episode for all agents. An important note to make is that the cultural parameter is also carried down from birth in the same fashion described in the Survival of the Fittest section above meaning there is also a chance for a mutation to occur here.

The goal of this chapter is to use both the Culture Parameter and Influence Rule in the creation of experiments which utilise our Artificial Society model in order to answer the following questions.

- Does a global culture form from the one-on-one interaction of agents? Section 6.2
- What is the impact of an individual on a society's belief system? Section 6.3
- Can we use the Artificial Society to model information propagation as if it were a connected graph? Section 6.4
- How does the environment affect information propagation? Section 6.5

¹A full visualisation of the cultural rule in use can be found at https://github.com/BrandenIngram/717431_Masters_Dissertation_ArtificialSocieties

6.2 The Emergence of a global Culture

Previously we mentioned that [Epstein and Axtell \[1996\]](#) observed a society which converged to a global uniform way of thinking. Convergence was a direct result of the interaction between agents. In this section, we demonstrate that even with our extensions to their implementation we can still generate a society which, given enough time, will converge. The importance of such an experiment is to analyse how individual interactions can affect global outcomes. For example, can a single individual lead to change on a global scale and if so how consistently.

As before in Chapter 5 we ran 50 simulations of 1000 episodes each and averaged the results. We ran the experiments using the Reproduction Rule with no control on births, Movement Rule 2, Regrowth Rule as well as the Influence Rule. The Culture parameter for our initial population was also randomly generated in a similar fashion to the other genes/parameters of the agent. In order to determine if convergence has occurred, we took two approaches; an analytic approach as well as a visual approach. For the analytic approach, we recorded the number of times two agents had differing viewpoints (A had a differing value in their culture parameter at position i to that of B) per episode. How this value changes over time can be seen in Figure 6.1 below.




Figure 6.1: Total number of cultural differences for the simple model

We can conclude that what emerges is a society that slowly converges towards a global uniform belief system and then due to a mutation an agent is born with a slightly differing belief which diverges the society and the process then repeats. This can be seen in Figure 6.1 where the number of differences initially is very high and slowly tends to decrease. It then spikes up again and once again tends to decrease, generating a cycle of convergence followed by divergence. It is also important to note that due to the averaging over 50 simulations we encounter noise as the behaviour does not necessarily happen at the same time across the simulations. In order to make the cyclic behaviour easier to observe we added an additional line which is the average rolling window. This line was generated by averaging the number of cultural differences of the 10 episodes either side of the current episode. From this we get a better sense of the up and down nature of the experiment.




Figure 6.2: Total number of cultural differences where mutation was prevented for the simple model

To truly conclude that the birth of an agent with a mutation in their culture parameter was to blame for this divergence we decided to run the same experiment, except this time there was no chance of a mutation in the Culture parameter. Figure 6.2 above is the resulting number of differences in this experiment. From this figure, we see that the number of differences ultimately converges to 0. Thus we can conclude that the society which would emerge where no mutation can occur or where no outside sources could influence an agent's beliefs would ultimately converge to a uniform belief system.

In order to get a better visualisation of this process, we consider a single simulation. The figures below depict the current state of the world at various stages of the simulation. To do this we converted the Culture parameter into an RGB colour whereby the first 3 elements are used to generate our red channel, the second 3 our green channel and the last 3 our blue channel.

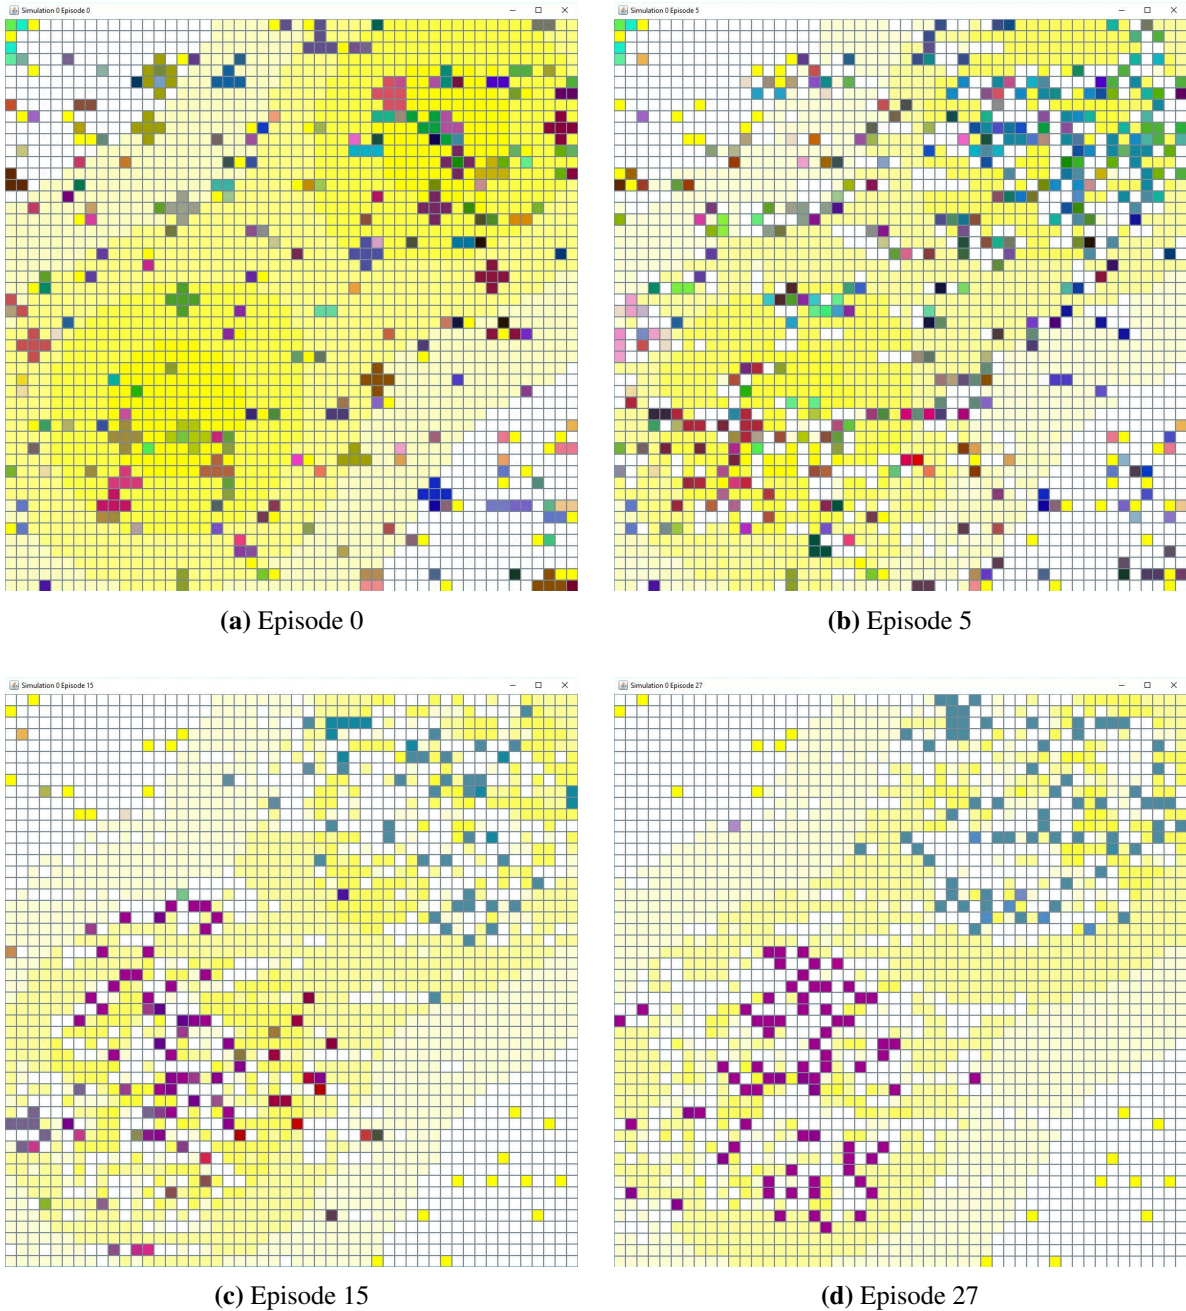


Figure 6.3: Visualisation of the Simple Model using the Influence Rule, Episode 0, 5, 15 and 27


From Figure 6.3a we conclude that our initial population starts off with a very diverse range of beliefs and this is due to our initialisation process. This conclusion can be seen by the wide range of colours across the world.

By episode 5 in Figure 6.3b we can already see small groupings of like colours if not similar colours. This can especially be seen in the top right region where we have a lot of blue or green agents. This shows our dialogue between agents is beginning to spread knowledge or beliefs throughout the society. We can surmise that two possibilities have occurred; either a single agent has convinced multiple agents to change their beliefs to that of his/her own or, and what is more likely, agents have entered this dialogue and both changed to some degree. This small change then permeates to more agents and leads to further


small changes occurring eventually resulting in what we see. Another reason why the second case is more likely is that an agent can only influence those within their vision, and even those within their vision have an equal chance of imparting their own influence on the current agent.

In Figure 6.3c the simulation has jumped further ahead to episode 15 and what we see is that roughly two factions have formed in the regions where the richest yield of resources are located. The population is visibly smaller which is consistent with our other results and this is due to our agents either being spatially located in regions extremely scarce in resources or the agent has poor genetics. Both of which would result in eventual death as seen previously. What can be concluded at this early stage of the simulation is that if some factor prevents interaction between people such as in this case we have a spatial separation, what emerges is two distinct ways of thinking or belief systems. It is important to note that both groups still have individuals who differ slightly from the local belief system. These individuals still have the possibility of influencing the overall belief system. How the environment affects the population's culture will be explored in greater detail in a later section.


By Figure 6.3d we have reached a situation where two distinct groups have formed where each member of each group has the same belief system. This is seen in the fact that all agents in the south-west have a pink colour whereas everyone in the North-East has a blue colour. Here we can conclude that in a local sense we have a society that has converged to a uniform belief system.




(a) Episode 35



(b) Episode 42



(c) Episode 105



(d) Episode 130

Figure 6.4: Visualisation of the Simple Model using the Influence Rule, Episode 35, 42, 105 and 130


In Figure 6.4a we observe that the populations of both groups are growing larger and what ends up happening is that as the population increases, the lack of resources forces agents to spread out and therefore the two groups end up meeting in the middle. In the centre, we then see that agents from the differing groups end up influencing each other. These influences permeate back towards both groups' respective centres. This means that even though these two groups only meet in the middle, the dialogue that occurs in the middle affects the entire population. This can be seen on the border where we observe agents whose colour is that of neither groups.

By episode 42 in Figure 6.4b we observe the effect of the process that started in Figure 6.4a. We no longer have two distinct groups but rather a number of new groups have begun forming, those being


grey, pink, dark purple and a light purple. The society can be considered to be in a state of cultural diversity which is constantly changing. This society is akin to two distinct belief systems trying to integrate into one uniform way of thinking.

The process of the influence rule continues even longer and what emerges is that the society eventually converges to a single uniform way of thinking. This can be seen in Figure 6.4c where each agent has pink-purple shade representing their beliefs.


We note that from Figure 6.4d that this stable society is not immune to change due to the possibility of a mutation occurring when an agent's beliefs are not carried down by either of their parents. This can be seen in the upper part of the world where a brown agent has been born. This agent has, in turn, managed to influence a couple of other agents. These agents with slightly different beliefs have the opportunity to influence the overall belief system. However, there is also an opportunity for the rest of the society to convince them of their populist beliefs. The outcome of this can be seen in Figure 6.5 below.



(a) Episode 135



(b) Episode 516



(c) Episode 996

Figure 6.5: Visualisation of the Simple Model using the Influence Rule, Episode 135, 416 and 996

We observe in Figure 6.5a that the second case described above ends up occurring. In this case, we see the world returns to its original uniform state of belief. This case has a far greater opportunity of occurring due to the fact that there are far more agents within the divergent agent's vision that have the uniform belief, than those who don't. This means there is a higher chance that the divergent agents will be convinced to conform to that of the pre-existing uniform belief system of the society.

In the Figure 6.5b above we have a similar case as to what was described previously where a new agent is born with a slightly different belief and once again said agent has the opportunity to influence the global belief system. However, where before the global system wasn't affected, here we see the other case in which an individual divergent agent or a group of divergent agents end up causing a great change

in the society. This does not mean that the society that emerges is one that follows the original beliefs of these initially divergent agents. What it means is that these divergent agents end up shaking up the world by introducing new ideas which lead to other ideas being generated. This is why you can see a notably chaotic society forming where there are many different colours representing different beliefs or ideas. The simulation then continues for a long period of time in this state of chaos where everyone is spreading slightly different ideas, and sharing what they believe. However, as we saw before eventually the society that emerges is one of a uniform belief system and this time we can see that it is different from that original uniform belief system. This result is depicted in Figure 6.5c. This is because we can now see that the society is now all relatively lime green compared to the purple of before.

From this simple influence rule which describes how an agent shares ideas with its neighbours, has emerged a complex system which evolves a uniform way of thinking from an initially very diverse population. Through both the analytic approach of calculating the number of cultural differences as well as our visualisation of the changing society, we have shown that our model, like that of [Epstein and Axtell \[1996\]](#), can converge to a uniform global belief system. In our approach, however, there is the opportunity for new beliefs to arise through the advent of a mutation in the Culture parameter. We have also shown that this event may or may not lead to a change on a global scale for the population. Finally, we have also demonstrated that isolated groups of individuals converge to their own local belief system. This local uniformity is only interrupted by the advent of coming into contact with an external belief system. These local groups also converged relatively quickly allowing us to conclude that the rate at which a belief system converges is proportional to the population size. This last observation was supported by the fact that it only took 27 episodes for two distinct groups to converge to a uniform belief seen in Figure 6.3d, whereas after the two populations meet it takes a further 70 episodes to converge as seen in Figures 6.4a and 6.4c.

6.3 Measuring the influence of individual agents

In this section, we study the effect of an individual agent on an entire population of agents by utilising the influence rule and cultural parameter once again. Measuring influence has been used in determining cluster stability [[Cerioli 1999](#)] for different clustering algorithms such as k-means clustering. K-means Clustering is a technique of grouping data points into clusters, where the number of groups is determined by the value k . In terms of their research, the aim of measuring a single data point's influence was to determine which data points had the biggest influence on the resulting clusters. The same idea can be applied to our model. Here our data points are represented by the agents and the actual value of the data points can be represented as our cultural parameter. We have drawn ties between our existing model and that of work previously done within this topic of measuring influence, however, we have not described the reason why it is important.

From the time we are born we are bombarded by external influences from parents, to teachers, to media, to friends to even the weather. Everything and everyone we interact with is an influence on oneself and our behaviour or beliefs are affected by this interaction. For example, a teacher presents you with an explanation of a new concept such as gravity. This explanation is suitably convincing and now you as an individual believe this concept to be true. Going forward if prompted for discussion on gravity you would present a similar case as your teacher did. Based on the outcome of that discussion and subsequent discussions, your beliefs would shift. This moulding of ideas is what we have already demonstrated previously whereby an agent's colour changes through interaction. There are two important ideas we touch on in this example. Firstly by being influenced by your teacher you are subsequently being influenced by everyone who influenced him/her. This chain can be followed all the way back to Sir Isaac Newton


and then even further to all the people who influenced him. It is by tracing the route further and further back, that we get a real sense of how powerful the influence of an individual can be on a society, whether their impact was good or bad. The second idea is that by knowing our own influences we can get a sense of our own biases, as well as the biases of the people we are interacting with. By understanding who these influencers are we can make more informed decisions.

In order to measure the influence of a single agent, we set up an experiment as follows. A “special” agent would be born at episode 50 and be placed at row 40 and column 10. All other agents were prevented from visiting this tile throughout the simulation in order to ensure that the tile was vacant. This “special” agent was displayed with a white border so that it would be easily seen in the visualisation. This agent was also defined with good genetic traits, for metabolism a value of 1 and for vision 5, as well as given a large amount of sugar at birth. This was to ensure that this agent survived until death by old age. This agent was also prevented from reproducing as it’s high sugar value would cause problems whereby really rich children would be born. This potentially could have led to the early death of the “special” agent. In order to incorporate this new agent, the influence rule was modified. Now if an agent entered a dialogue with this “special” agent they would (based on a percentage chance) match the beliefs of the “special” agent. This was done by changing their cultural parameter to that of the “special” agent. The influence rule remained the same as in previous experiments, for when neither of the agents are a “special” agent. Mutation of the cultural parameter was also prevented in order to ensure the population converged to a uniform belief system. Once this convergence had occurred the uniform culture was recorded in order to compare it to that of the “special” agent. The culture parameter of the “special” agent was defined as all zeroes as depicted below.


index	0	1	2	3	4	5	6	7	8
value	0	0	0	0	0	0	0	0	0

Table 6.2: “special” agent’s cultural parameter


Figure 6.6 below depicts the visualisation of the experiment outlined above whereby there was a 10% chance that a normal agent would take on the “special” agent’s beliefs. The same episodes have been displayed in all three of the experiments to come in order to maintain a consistency in the visualisations. This was also done to get a better sense of the differences between each set at the same point in time.




(a) Episode 50




(b) Episode 64



(c) Episode 111



(d) Episode 130




(e) Episode 350

Figure 6.6: 10% chance of Information propagation with “special” agent


In Figure 6.6a we see the birth of the “special” agent depicted in the bottom left with the white border and black fill. In Figure 6.6b we can see that the “special” agent has influenced some but not all of the agents within its vision and this is due to the 10% chance to influence. We can also note that the rest of the agents not influenced by the “special” agent is beginning to form a uniform belief, this can be seen by the majority of the agents being shades of red and pink. By Figure 6.6c we can see that the “special” agent’s influence has spread further, however, it has been diluted or tainted by the other agents in the region. This is seen by a large number of brown agents who share a large portion of their beliefs with the “special” agent, however, not all their beliefs are the same. We note that in Figure 6.6d the “special” agent has died of old age. Since this agent would be 80 years old it has a 33% chance of dying. This is an important observation as we can now see the extent of its influence during its lifespan. We can see that it has only managed to influence the bottom left portion of the world, however, the “special” agent’s beliefs live on in those that it influenced, so there is still the opportunity of this belief system spreading. This is exactly what is seen in Figure 6.6e. Here it is seen that the world has converged to be completely black. This result would not happen every single time as we could see that the red faction in the previous figure, could have become the predominant belief system.

This experiment has shown us that individual agents can have a big influence on the global population. It has also shown that “special” agents symbolic of powerful individuals in the real world can have a massive role in shaping the way the majority thinks and that an individual can have a lasting effect even after their death. This is because although that person may die, their teachings and ideas live on in those people that they influenced. These people, in turn, have the power to influence those around them, which means an individual’s influence can be seen in generations that came long after themselves. In conclusion, we have shown how once again using a simple rule we can model how an individual’s influence affects a large population over a period of many generations.


Figure 6.7 below depicts the visualisation of the experiment outlined above whereby there was a 0.1% chance that a normal agent would take on the “special” agent’s beliefs. This experiment was conducted in order to determine if we reduce the chance the “special” agent has of spreading their influence, does this, in turn, affect the reach of this agent overall. Reach is the number of people one can influence whether by direct contact or by indirect contact. This indirect contact occurs when an agent interacts with someone who has been influenced previously by the “special” agent.




(a) Episode 50




(b) Episode 64



(c) Episode 111



(d) Episode 130



(e) Episode 350

Figure 6.7: 0.1% chance of Information propagation with “special” agent


Once again in Figure 6.7a we see the birth of the “special” agent and we note that the world is currently predominately green. However, there is by no means a uniform culture at this point. In Figure 6.7b we see that more of the population has become green. Secondly, the “special” agent has only influenced a small portion of those around it. When comparing Figure 6.7c with the figure in the previous experiment Figure 6.6c at this same time, we observe that with a lower percentage chance of 0.1% we also get fewer people who were influenced. This observation is further validated in Figure 6.7d where there are only a few agents with black or brown colouring. It is also important to note that the “special” agent has died off at this point. In contrast to the previous experiment, we see that in the final Figure 6.7e the population has converged to a green colour which is different from that of the “special” agent. The converged cultural parameter of the agents in this last figure is depicted in the table below.

index	0	1	2	3	4	5	6	7	8
value	0	0	0	0	1	2	0	0	0


Table 6.3: “special” agents’ cultural parameter

As can be seen above there are differences, but there are also similarities. These similarities would have arisen from the indirect reach the “special” agent had. They can also be the result of random chance as initially we see a predominantly green society which was what emerged. This green culture originally shared some characteristics of the “special” agent and was not a result of an interaction. This experiment shows that the reduced ability of the “special” agent as a result of the lower percentage resulted in a society which didn’t share all of its’ beliefs with the “special” agent which is in contrast to the previous experiment.

Figure 6.8 below depicts the final set of visualisations of the experiment outlined above whereby there was a 0.1% chance that a normal agent would take on the “special” agent’s beliefs. The “special” agent was also defined with a vision of 10. This experiment was designed in order to analyse the effect of reach. We hypothesised that if we increase the “special” agent’s vision we would, in turn, increase their reach, resulting in more people being influenced by that single agent.




(a) Episode 50




(b) Episode 64



(c) Episode 111



(d) Episode 130



(e) Episode 350

Figure 6.8: 0.1% chance of Information propagation with “special” agent and vision 10

In Figure 6.8a we once again see the birth of the “special” agent into a society with a predominately pink and purple belief system. Already by episode 64 in Figure 6.8b we can see that the “special” agent has influenced a number of agents. This is more than the previous experiments at this same point. By Figure 6.8c we see that almost half of the population has now been influenced. This is far more than the previous experiments. In this experiment the “special” agent took 3 additional years to die and can be seen in Figure 6.8d. We also note that the majority of the bottom left quadrant share the beliefs, if not very similar beliefs, to that of the “special” agent. As in the second experiment, our population doesn’t converge exactly to the belief system of the “special” agent, however, it is very similar as seen in the cultural parameter depicted below.

index	0	1	2	3	4	5	6	7	8
value	0	2	0	0	0	0	0	0	0

Table 6.4: “special” agents’ cultural parameter

Here we can see it only varies at one position. Once again it is important to note this result might not occur exactly the same way due to the probabilistic nature of the interactions between agents. From this experiment, we saw that whilst the “special” agent was alive it was able to reach more agents than that of our previous experiments due to the increased vision even though it still had a really low chance of spreading influence.

Table 6.5 below was generated to get a better understanding of how the behaviours observed in each experiment compared to each other.

Experiment	Episodes to Converge	Agents Influenced at Death	Converge to “special” agent
10% chance of propagation	302	67	Yes
0.1% chance of propagation	276	51	No
0.1% chance of propagation and Vision = 10	290	67	No

Table 6.5: Results of “special” agent experiments

From Table 6.5 We observe that the lower the chance to propagate results in a lower number of people who have been influenced by the time the “special” agent dies. We also observe that when increased the vision of the “special” agent that this change counteracted the effect of the lower chance of propagation. This is seen in the third experiment of Table 6.5 where the number of agents influenced is higher than the second experiment. We note that the time till convergence is highest for the experiment with the highest chance of propagation, this might be due to the larger impact of the “special” agent which is disrupting the natural convergence process of the world. This is also the only experiment which converged to the cultural parameter of the “special” agent which further leads us to believe that the slow convergence is because of the “special” agent and not a random happening.

In all of the experiments above we have shown how impactful an individual agent can be on the global belief system of a society as well as how quickly this individual’s influence can spread. We have also measured this impact by comparing the Culture parameter of the “special” agent with that of the global culture which emerges. In all experiments, there were a majority number of similarities between them. Another important observation that was made is that in all cases the society did not reach a converged belief system within the “special” agents lifetime. This means that in order for the world to have converged to that of the “special” agent required those who were influenced to continue spreading those beliefs to those around them. This once again goes to show how an individual can still have lasting effects on a population well after their own death. The final conclusion to be made is that an individuals influence


is greatly dependant on their reach. Reach in turn being greatly affected by the number of people they can directly influence as well as how effective they are at conveying their message. This was seen in the experiments using differing visions along with different percentage chances to influence.

6.4 Cultural parameter as a model of information propagation


The way in which information spreads within a society has changed drastically in just a short period of time. Previously where messages were transported physically from person to person in many forms such as mail or telegraph, now as a result of many technological advancements such as the internet we can send messages instantly around the world. However, the fundamental idea of information spread remains the same. This idea is that information is spread in the form of a network where people are the propagators of information and the channel they use is irrelevant simply just forming a link between people. If we consider this type of network we can draw a parallel with our model which utilises the cultural parameter. Consider the cultural parameter as the information an agent has. Now when an agent interacts with another agent as previously described above, we can consider this as a propagation of information from one agent to another. This is very similar to what we would consider a word of mouth model where people spread information to those within their circle and in turn, those people spread the information to others in their own social circles. However, we can interpret these agents as any type of player in a network such as a company or a website. In this section, we demonstrate how our Artificial Society can be used as a new way of modelling information propagation as well as what aspects of an individual affect information propagation.

In order to demonstrate how the model could be used to model information propagation, an experiment was designed. For this experiment, the initial population of agents were all given the same culture parameter values. Once again a “special” agent would be born at episode 50. The design of this “special” agent is the same as the experiment prior. Agents were only allowed to interact with this “special” agent and when they did they would take on the “special” agents cultural parameter as their own. They would in turn also become “special” and therefore this would allow them to in turn to propagate the recently acquired information. Finally, due to the chance of a mutation occurring from reproduction, there is a chance of additional information being spread. This can be considered misinformation or an occurrence whereby only some of the information gets across to the other person.


For the first experiment, we allowed for a 10% chance for information to propagate from “special” agents to normal agents. The visualisation of this experiment is shown below.




(a) Episode 49




(b) Episode 50




(c) Episode 51



(d) Episode 52



(e) Episode 54




(f) Episode 57


Figure 6.9: 10% chance of Information propagation

In Figure 6.9a we can see that the population of agents evolves in the same fashion as seen in previous experiments. We can also see the empty yellow tile in the bottom left where our initial "special" agent will be born. In Figure 6.9b we see the birth of our "special" agent. Figure 6.9c shows how information is beginning to spread amongst the nearby agents. It is important to note that the influence rule is still working the same, whereby agents only spread influence with those within their vision. The 10% chance of information spread also means that not all agents within vision will receive the information. By Figure 6.9d we can already see a large portion of the population has already received the information, showing that even with a 10% chance to spread information it still propagates really quickly. By Figure 6.9e the information has reached almost the entire population. Finally in Figure 6.9f information has propagated through the entire population in just 7 episodes. An interesting observation from this set of figures is that there was no occurrence of misinformation or if it did occur it was very quickly corrected. This is again due to the high rate at which information is spread thus reducing the impact of any occurrences of misinformation.


In Figure 6.10 depicted below we ran the same experiment, however, this time we drastically reduced the chance to spread information to 0.1%.




(a) Episode 50




(b) Episode 54



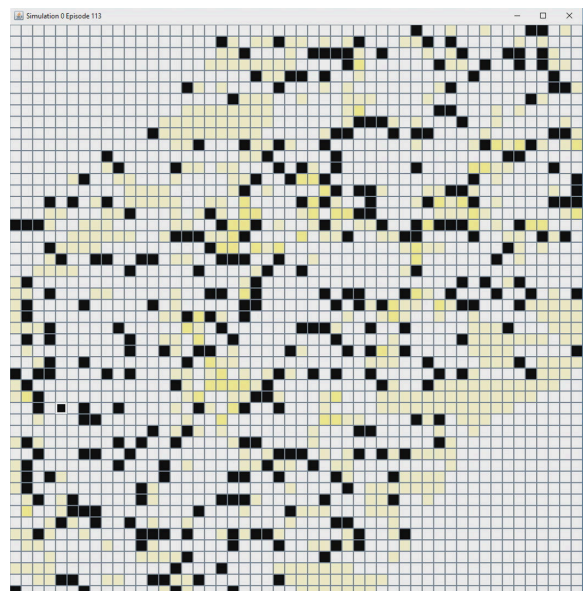
(c) Episode 61



(d) Episode 77



(e) Episode 97



(f) Episode 113

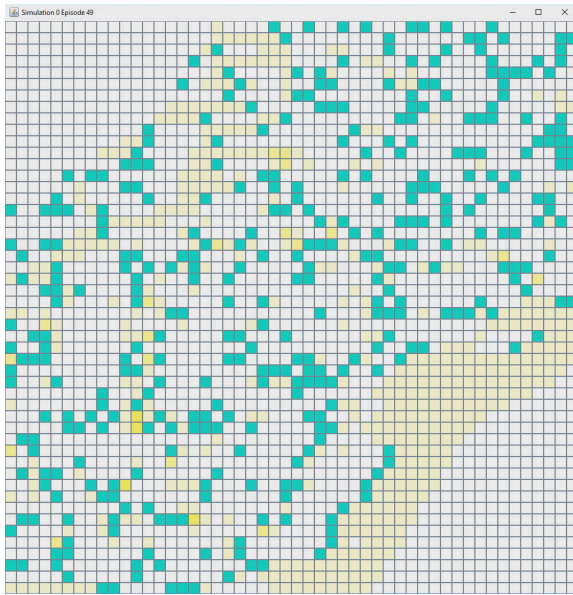
Figure 6.10: 0.1% chance for Information propagation

Again in Figure 6.10a we see that the population that evolved was relatively the same as that in all previous experiments. This is important as we want to keep our initial world that our “special” agent is born into relatively the same. This is because we are only interested in the effect of changing the chance to spread information. By Figure 6.10b there is still no information spread to any other agents. This is a great contrast from that of the previous, where half the population had already received the information at this point. In Figure 6.10c finally information is beginning to spread, however, we also note other colours namely pink and dark blue. This is a result of the mutation occurring, as previously stated, this represents a spread of misinformation. In Figure 6.10d we can see that the rate of spread is increasing and this is due to the fact that as more agents have the information, there are more opportunities to spread the information. This is, however, restricted due to the size of the population and the fact that agents without the information will be found only on the borders, but as a general rule, it holds true. By Figure 6.10e almost the whole population has received the information, however, there is still a notable number of agents who obtained misinformation. Finally in Figure 6.10f we no longer see any agents with any misinformation and the information has propagated throughout the whole population. This took 56 additional episodes when compared to the previous experiment.

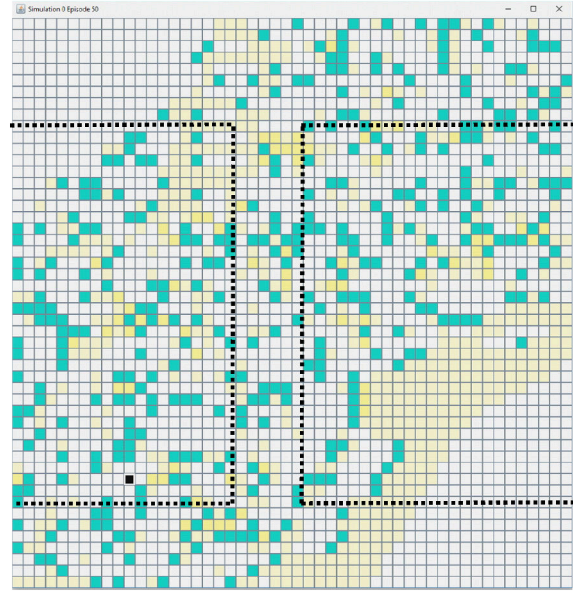
From these two experiments, we have shown how we can use our model with very few modifications in order to demonstrate how information can propagate throughout a population through the interactions at an individual level. We have also shown by adjusting our chance to spread information we can dramatically affect the time in which information takes to propagate throughout the entire population. An important note on the density of the population can also be made, where large populations of dense individuals can propagate information far quicker than that of a sparsely populated society. This is because there is a greatly increased number of interactions in a dense population as many more agents fall within each other’s visions. The converse is also true for small populations and taken to its extremes we could see no propagation due to isolated individuals or groups being too far to interact. This can be seen throughout history in general, where very small societies have developed slower. A modern example of this are the primitive tribes of the Andaman and Nicobar islands [Gautam 2005] who operate as traditional hunter-gatherers and serve as a study of the antiquity of human evolution. Their isolation from the outside world, as well as their willingness to remain isolated, has had a massive impact on their development as a society. Finally, we could also change the agents’ vision parameter which would, in turn, affect the rate at which information propagates.

6.5 Cultural parameter as a model of information propagation based on topography

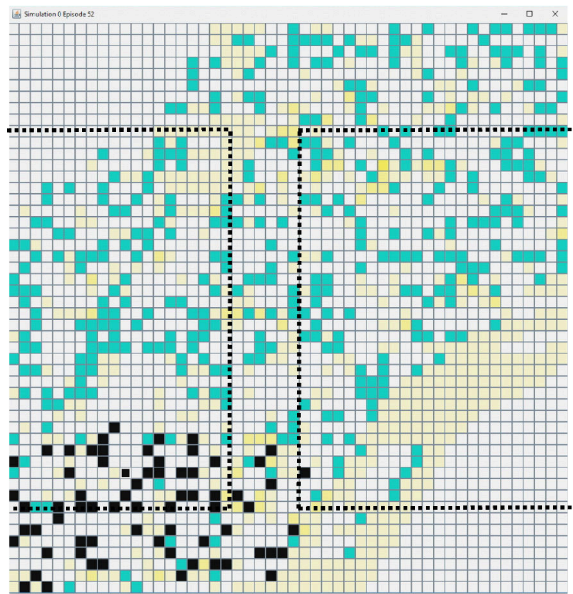
In the previous experiment, we only consider the agents and how they affected the spread of information within the population, as well as how important population size is. Another important aspect to consider is the topography of the environment in which your network of agents exists. As mentioned above the primitive tribes were located on islands, these islands and the fact they are remote, present a natural barrier to the spread of knowledge. Thus increasing, even more, the level of isolation of these primitive societies. It for this reason, that it is important when modelling the propagation of information to take into consideration the environment, that the network of people exists within. Therefore, the first goal of this section is to demonstrate how we can modify the model used in the experiment prior in order to take the environment into account. The second goal is to measure the impact an environment can have on the propagation utilising different percentage chances of spreading influence in the same environment.




(a) Episode 49




(b) Episode 50




(c) Episode 52



(d) Episode 57



(e) Episode 62




(f) Episode 100

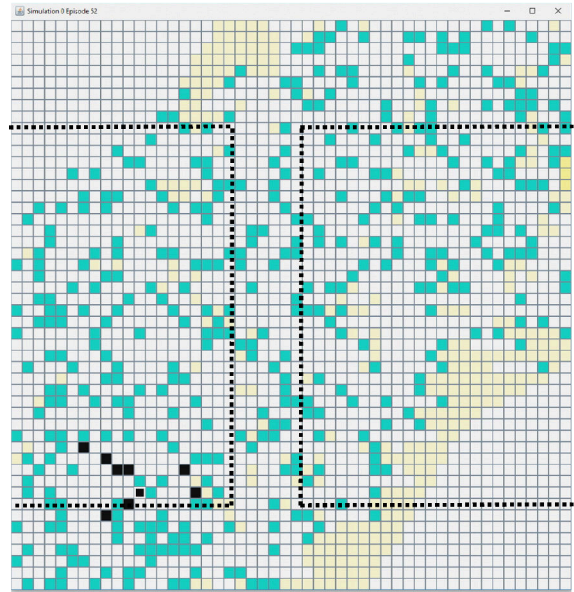
Figure 6.12: Long Bridge, 10% chance experiment

In Figure 6.12a we see the same initial world arrangement as before. In Figure 6.12b we again see the birth of the initial “special” agent. By Figure 6.12c we can see that information has begun to be spread and we observe that the information has a preference to spread horizontally rather than vertically which is accordance to that of the “long bridge” world we designed. The topology is representing a natural or artificial (depending on the analogy being used) barrier to the spread of information. In an artificial sense, we can consider this barrier to be a network where firewalls prevent information from reaching users. However, we will consider the natural case for these experiments where the environment represents a bridge crossing a gorge. In Figure 6.12d this observation continues to hold true as information is being spread up along the narrow bridge in the middle while it is still being prevented from spreading in the left and right regions. By Figure 6.12e the information has reached the top of the world. Those agents that are present in the middle left and right have not received the information and those everywhere else have. It is important to note that we do observe agents which have received the information who are situated on regions who fall under the “9” section of the topology. This is due to the way in which the rule has been implemented whereby we consider if the middle point between two agents falls on a “9” or not. This therefore allows for agents who are situated on a “9” tile to receive info they interacting with another agent who results in their middle tile not falling on the “9” tile.


Figure 6.13 and Figure 6.14 to follow represent the experiment using the “long bridge” with a 1% and 0.1% chance for information to propagate respectively. Discussion and analyses of both these experiments appear after Figure 6.14




(a) Episode 50




(b) Episode 52



(c) Episode 57

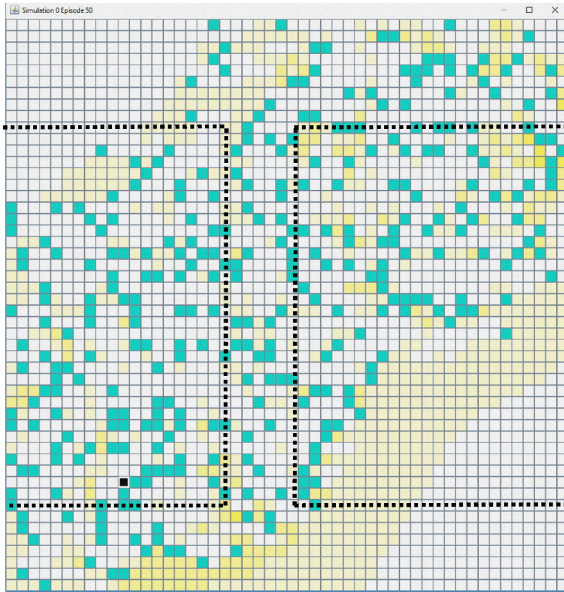


(d) Episode 62

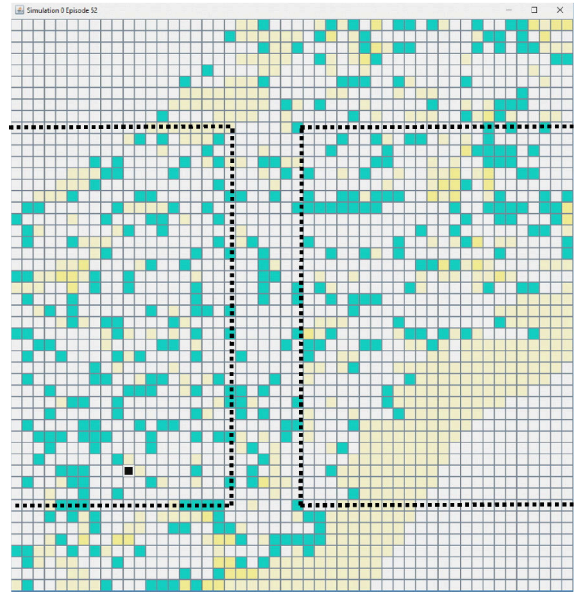


(e) Episode 100


Figure 6.13: Long Bridge, 1% chance experiment.




(a) Episode 50




(b) Episode 52




(c) Episode 57



(d) Episode 62



(e) Episode 100



(f) Episode 167

Figure 6.14: Long Bridge, 0.1% chance experiment.

The purpose of the previous two experiments and their corresponding sets of figures was to show that the original experiment still holds true, whereby if we reduce the chance of information spread this, in turn, increases the opportunity for misinformation as well as increasing the time taken to spread information. Table 6.6 shows the number of episodes required for information to reach the top of the world for each of the three experiments.

Percentage chance to spread	Number of episodes
10%	12
1%	49
0.1%	117

Table 6.6: Long bridge experiments comparison

From Table 6.6 we conclude that the greater the percentage chance of information spread the quicker it takes for information to spread within a society. This is also in accordance with the previous experiment conducted which validates the result. More specifically if we decrease the percentage chance by a factor of 10. However, we see an unequal rate of increase in the number of episodes. Firstly it increases by a factor of 4 and then by roughly 2.

In Figure 6.14f it is important to note that it appears that information has spread to the left region. This is not due to the usage of the influence rule, but rather that of reproduction as the information can still be carried down to children, just as the before and since this experiment goes on for so long before reaching the top, a large number of children were born in this region resulting in the appearance of information spread into this region.

Additional experiments were run using a wider bridge as the information restricting topology, however, the behaviour observed was the same as the experiments already conducted. The one difference being that agents which would have been on the outskirts were now able to receive the information. The general behaviour of information spreading right first to the bridge, then up through the bridge then right again was still present.

These experiments show that through this simple implementation we can generate a society where information spread is not only dependant on the interaction between agents but also on the topography of the environment they exist in. We have also validated previous results by demonstrating that even when we are considering the environment, the percentage chance of agents to spread influence has a direct effect on the rate of propagation.

6.6 Conclusion

In this chapter, we set out to determine if a global culture would arise through the one-on-one interaction. In terms of this question we showed that, even with our additional modifications to the basic implementation proposed by Epstein and Axtell [1996], we can, in fact, generate a society which converges to a uniform belief system. We also discovered that through gene mutation we can simulate the effect of introducing outside influences to a population. Secondly, we set out to measure the impactfulness of individuals on the overall belief of a society. Through our experimentation, we were able to determine not only the degree to which a particular individual has an effect on the society but also what traits lead to greater influence. We also observed that in all cases it was those who the “special” agent had influenced who would ultimately continue spreading the belief well after the “special” agent had died.

Thirdly we wanted to demonstrate the application of our model for information propagation. In answer to this question, we demonstrated how our agents can spread information within their vision. We also showed the aspects which affect the rate of information propagation as well as how an environment can play a role.

One limitation of the implementation described above is that all agents have a uniform acceptance of what they learn from other agents based on a percentage that all agents abide by. For example A compares tag with B, their tags are different, therefore B accepts the tag of A based on a fixed percentage. This is unrealistic as people are not all equally susceptible to changing their beliefs. Some people are easily influenced while some others are more stubborn.

In conclusion, the Influence Rule along with the Culture parameter can be used in a number of interesting ways, to model a variety of problems all of which reside within the framework of this artificial society model. The complex behaviours here have all arisen from simple rules and parameters contained within this multi-agent system. Further use of this rule can be used in determining outcomes based on preference. For example, reproduction can be implemented so that agents determine possible partners based on their cultural beliefs or preferences. An additional example would be in terms of trade whereby agents only trade with other agents of similar beliefs. This second example will be expanded and implemented in the following section.

Chapter 7

The Trade Model

7.1 Introduction

Trade can be described as the willing exchange of goods or services between two or more actors. One could also consider trade, along with the foundation of trade markets, to be a very important innovation in human history and the leading proponent of development [Davies and Quinlivan 2006]. A trade market being an instrument which has allowed trade on a global scale. A market is a place where buyers and sellers come together and engage in buying and selling of goods and services. Such a market can be a physical place, or it can be an electronic infrastructure created by various financial institutions. This can be seen by the fact no business could compete with others if it were to produce everything it required to operate in-house as it would not be affordable. Andy George [2018] created a video series called “How to Make Everything” demonstrating what it would cost if we were to make a sandwich from scratch. They discovered it would cost around \$1500 and take 6 months, they compared this to around \$15 and half an hour if they were to purchase the ingredients in a store. This large saving can be contributed to voluntary trading and is the reason we can afford goods and services. It is, therefore, the goal of this chapter to implement a rule which allows agents to trade amongst themselves, in order to analyse the effect of trade. The second goal is to see whether the results obtained by Epstein and Axtell [1996] can be extended to a three resource environment. Lastly, this section analyses the effects of different types of trades as well different trading strategies of individuals.

The inspiration for the Trade rule and how it works was drawn from Epstein and Axtell [1996], where they simulated a two resource trade environment. We expand on this idea by scaling up to a third resource as well as by introducing trading strategies and a separate experiment where cultural differences impact on who agents are willing to trade with. Similar to when we increased the size of the cultural parameter, we again chose to add the third resource to determine if the results obtained by Epstein and Axtell [1996] could be replicated in a more complex environment. All of this is explained in further detail in this chapter. The trade rule operates in a similar fashion to that of the cultural rule. The trade rule describes how two agents enter into a negotiation in order to hopefully trade resources between each other for a mutual benefit. In order to simulate a trading environment, two additional resources are added to the environment of our simple model, namely spice and food. The way spice was distributed was similar to that of sugar, however, a rotation was applied. This means that overall spice occurs in the same amounts across the grid just in different locations. Food was distributed completely differently, however, whereby the food capacity of a tile was randomly generated as a value in $\{0, 1, 2, 3, 4\}$. The initial endowment of sugar, spice and food were each randomly initialised on the range $[5, 45]$. The reason for the increase was due to the harsher requirements of the environment. This difficulty is demonstrated in the results of

the experiments to follow.




Figure 7.1: Levels of sugar across the world. Darker yellows indicate higher yields




Figure 7.2: Levels of spice across the world. Darker yellows indicate higher yields




Figure 7.3: Levels of food across the world. Darker yellows indicate higher yields

Figures 7.1, 7.2 and 7.3 depict the distribution of yields for each of the three resources. Once again the darkest shade of yellow represents a value of 4 all the way down to white having a value of 0.

The corresponding metabolisms were also added to the agents as internal parameters for each of the new resources, both of which would be necessary for survival. This means that if an agent runs out of any of the three resources, they will die. The movement rule was also adjusted in order to incorporate the additional resources. Now instead of seeking the highest yielding sugar tile, an agent must search for the tile which offers the greatest net benefit. We defined a welfare equation inspired by the one derived in Epstein and Axtell [1996] for their two resource environment in order to accomplish this task and it works as follows:

$$\text{Welfare}(r_1, r_2, r_3, su, sp, fo) = (r_1 + su)^{\frac{m_1}{\sum_{i=1}^3 m_i}} \times (r_2 + sp)^{\frac{m_2}{\sum_{i=1}^3 m_i}} \times (r_3 + fo)^{\frac{m_3}{\sum_{i=1}^3 m_i}} \quad (7.1)$$

where

r_1 represents the amount of sugar an agent currently has.

r_2 represents the amount of spice an agent currently has.

r_3 represents the amount of food an agent currently has.

su represents the amount of sugar that can be gained from moving to this location.

sp represents the amount of spice that can be gained from moving to this location.

fo represents the amount of food that can be gained from moving to this location.

m_1 represents the sugar metabolism of an agent.

m_2 represents the spice metabolism of an agent.

m_3 represents the food metabolism of an agent.

$$\text{Welfare}(r_1, r_2, su, sp) = (r_1 + su)^{\frac{m_1}{m_1+m_2}} \times (r_2 + sp)^{\frac{m_2}{m_1+m_2}} \quad (7.2)$$

Now similar to the way it worked in the original movement rule, the agent looks out in their vision based on the parameter of the same name. Then, where before the agent would select the tile with the highest sugar, they now check the welfare they would gain from moving to each tile and pick the one

which results in the highest return (the tile which returns the highest value from the welfare function). The exponent of each component represents the relative need of that resource. For example $\frac{m_1}{\sum_{i=1}^3 m_i}$ represents the relative need for sugar. If m_1 is the highest metabolism when compared to food and spice the resulting exponent will be the highest relative to the others. Then the tiles which offer a higher su value will result in a higher overall value for the Welfare function. The opposite can be concluded if m_1 was relatively small, we would see the agent not favouring high yielding sugar tiles over equally high and in some cases lower levels of other resources. It is also possible to calculate an agent's current welfare using Equation 7.1 with an input of 0 for parameters su , sp and fo .

Lastly a way of calculating the Marginal Rate of substitution between two resources was required in order to determine trade direction and price. This equation was defined as follows:

$$\text{MRS}_{r_1 r_2} = \frac{\left(\frac{r_1}{m_1}\right)}{\left(\frac{r_2}{m_2}\right)} \quad (7.3)$$

where

- r_1 represents the amount of resource1 an agent currently has.
- r_2 represents the amount of resource2 an agent currently has.
- m_1 represents the metabolism of an agent w.r.t. resource1.
- m_2 represents the metabolism of an agent w.r.t. resource2.

This formula was also inspired by the derivation outlined by Epstein and Axtell [1996] from which they derived MRS as follows.

$$\text{MRS} = \frac{dw_2}{dw_1} = \frac{\frac{\partial W(w_1, w_2, 0, 0)}{\partial w_1}}{\frac{\partial W(w_1, w_2, 0, 0)}{\partial w_2}} = \frac{\frac{m_1}{m_t} w_1^{(m_1 - m_t)/m_t} w_2^{m_2/m_t}}{\frac{m_2}{m_t} w_1^{(m_2 - m_t)/m_t} w_2^{m_1/m_t}} = \frac{m_1 w_2}{m_2 w_1} = \frac{\frac{w_2}{m_2}}{\frac{w_1}{m_1}} = \frac{\tau_2}{\tau_1} \quad (7.4)$$

where

- m_1 represents the metabolism of an agent w.r.t. resource1.
- m_2 represents the metabolism of an agent w.r.t. resource2.
- m_t represents the sum of both metabolisms.
- w_1 represents the wealth of an agent w.r.t. resource1.
- w_2 represents the wealth of an agent w.r.t. resource2.
- τ_1 represents the number of turns an agent can survive without receiving more of resource1.
- τ_2 represents the number of turns an agent can survive without receiving more of resource2.
- $W(w_1, w_2, 0, 0)$ represents the welfare equation as seen in Equation 7.2 with inputs $w_1, w_2, 0, 0$.

The reason for using Equation 7.4 with input values of 0 in the welfare equation, is that we are looking at the current welfare of the agent, not the projected one. We note from both Equations 7.3 and 7.4 the MRS of an agent is dependant on the genetics or biology of that agent since the metabolism of an agent drives its need for obtaining a particular resource. From the Equation 7.3 we can consider the MRS of resource 1 with respect to resource 2. This equation allows us to determine the relative need of one resource with respect to another. For example, let us consider the MRS of sugar with respect to spice. If the value is greater than 1 that means our numerator ($\frac{r_1}{m_1}$) is greater than our denominator

$(\frac{r_2}{m_2})$). This implies that we have a greater need for the resource represented by the denominator, in this case, we would have a greater need for spice relative to sugar. Therefore this agent would be willing to trade sugar for spice. Now if the MRS value is less than 1 that means our numerator ($\frac{r_1}{m_1}$) is less than our denominator ($\frac{r_2}{m_2}$). This implies that we have a greater need for the resource represented by the numerator, in this case, we would have a greater need for sugar relative to spice. Therefore this agent would be willing to trade spice for sugar. By comparing the MRS's for the same resources between two agents, we can determine which agent has a greater need for which resources. This allows us to determine the direction in which the trade occurs. For example, if A has a higher relative need for sugar over spice than B, then A will trade spice to B for sugar. ¹

A basic description of the Trade Rule works as follows. If A sees B within their vision they enter the negotiating process. The negotiation proceeds as follows:

- For each resource combination (e.g. sugar → spice):
 - The MRS was calculated for both agents as well as their current welfare.
 - The price is calculated to be the geometric mean of the MRS's, which can be seen in Equation 7.5
 - From the respective MRS's the direction of the trade is determined in the same fashion as outlined above.
 - The trade is simulated to determine if it is mutually beneficial. This is achieved by comparing both agent's welfare before and after the trade if it were to be conducted.
 - If it is mutually beneficial then the trade is actually conducted.
 - This cycle continues until it is no longer beneficial to trade.

The price or geometric mean is calculated as follows:

$$p(\text{MRS}_A, \text{MRS}_B) = \sqrt{\text{MRS}_A \text{MRS}_B} \quad (7.5)$$

where

MRS_A is the MRS of Agent A for two particular resources

MRS_B is the MRS of Agent B for two particular resources

7.2 The benefits of Trade

Now that we have generated an environment which is conducive to trade as well as a rule which governs how trading occurs between agents, we can run our simulation to determine if it is beneficial for the society. In this section, we compare different metrics (age, population and wealth) to determine, whether or not allowing agents to trade improves the performance of the society. To do this we ran two experiments both of which incorporated the new three resource environment. The agents also for both experiments

¹A full visualisation of the trade rule in use can be found at https://github.com/BrandenIngram/717431_Masters_Dissertation_ArtificialSocieties

needed to fulfil their individual needs for each of the three resources in order to survive. The only difference between experiments was the initial distribution of agents as per all previous experiments as well as in one experiment the agents could trade and in the other they were prevented.




Figure 7.4: Total population of trade vs no trade for Trade Model

From Figure 7.4 we can see that initially, the population reacts the same regardless of whether agents are allowed to trade or not. There is the initial population collapse that we have seen in all previous experiments. Without trades, the population does dip to lower levels than in previous chapters, and this is evidence of the environment being even more difficult to survive in. This difficulty has arisen from the introduction of the additional resources and the need for agents to obtain suitable amounts of each. As the simulation continues what emerges is the fact that allowing agents to trade has a positive effect on the population leading to a stable population size of around 200, double that of the population when agents were not allowed to trade. Thus we can conclude that allowing trade increases the carrying capacity of an environment. The carrying capacity being the maximum sustainable population size.




Figure 7.5: Average age of trade vs no trade for Trade Model

The population that emerges when agents are allowed to trade tends to have a higher average age than that of a society where agents are not. From this, we conclude that not only are agents surviving longer due to the increased average age but they also survive in greater numbers.




Figure 7.6: Average wealth of trade vs no trade for Trade Model

When considering Average Wealth the previous benefit to trade is no longer seen, however, this is due to the fact that you now have far more agents competing for the same quantity of resources. This means that in larger populations our average wealth is lower than that of a smaller population.

In conclusion, we have demonstrated that even though on an individual level, trade results in a lower average wealth, the society ends up supporting a much greater population than that in which no trade

occurs. The average wealth of the society under trade is also only slightly lower, however, it is still far greater than the minimum required for survival which, if an agent had the worst genetic traits would be 4. We have also demonstrated that we were able to generate a stable society in a similar fashion to that of Epstein and Axtell [1996] even with our additions to the environment and the Trade Rule.

7.3 Global trade versus local trade

Now that we have demonstrated that trade is beneficial to a society, we need to look deeper into the Trade Rule itself. The most obvious element which affects how an agent trades is their vision parameter. In the previous experiments with trade, the vision parameter meant that agents could only trade with those agents within their own vision. This can be considered a local trade. In order to simulate a global trade scheme or a trade market, we implemented an additional trading rule. This rule worked exactly as before except now an agent would attempt to trade with everyone each episode. This means that an agent can access the resources gathered by an agent on the other side of the world. This also has the potential for allowing agents the ability to survive in regions which were too inhospitable previously. Therefore the goal of this section is to investigate the effects of global trade.

The comparison between these two trading rules can be seen below.




Figure 7.7: Total population per trading strategies for Trade Model

When it comes to population size, in Figure 7.7 we can see that the global policy also performs vastly better than that of the no-trade policy, however, it follows a very similar trend to that of the local trade. We can conclude that when comparing the two trade rules there is no notable difference in the effect they have on population size. A case could be made stating that the global trade performs better by converging to a slightly higher stable population but in general the difference is not large. The slight increase in population could be due to the fact that in the few regions previously unsuitable for life (corners of the world), agents can now survive.




Figure 7.8: Average age per different trading strategies for Trade Model

When considering the average age graph Figure 7.8, no notable disparity can be seen from any of the schemes, other than the slightly older population for the local trade scheme.




Figure 7.9: Average wealth per trading strategies for Trade Model

A look at the average wealth graph Figure 7.9 reveals an interesting result. Here the global trade scheme resulted in a lower average wealth, this can partially be accredited to the slightly higher population. Where the global trade policy would become more beneficial would be in a world where populations didn't have access to resources within their own vision. For example, two regions separated by water whereby agents could not travel or interact with the other island members if it were not for a global trade scheme.

The analysis of our global trading scheme reveals many interesting insights, especially when compared to that of the original local trade rule. Most notable of those was the fact that it only resulted in a slight increase in population size. We conclude that given a more extreme environment this gap would increase further. This is because the environment would not serve as a natural barrier when it came to the trading of resources nor would an agent's vision play a role in preventing trade. In general global trade results in far more trades and, although this results in a lower average wealth, that average is still far above the minimum required for life. We can conclude that without altering how resources are collected or how reproduction occurs, that global trade results in our highest achievable stable population.

7.4 Naivety

When making a trade, an individual determines whether the outcome of following through with it results in a net benefit. If it does the agent accepts the trade or negotiates for a greater benefit. If it doesn't the agents either re-evaluate their perceptions of the trade or they renegotiate for a more suitable trade. However, people aren't perfect and many of us will have made a bad trade. This is not always out of ignorance but merely because we do not always have perfect information about the trade we have entered into. It was from this logic that a naivety factor was incorporated into the trading rule. The goal of this section is to analyse how making bad trades affects individuals and whether or not the society evolves differently based on this additional factor.

How the naivety factor was incorporated is as follows. We first recall that a trade would be carried out if the trade increased the welfare of both agents in question. This was achieved by checking the welfare before the trade with the resulting welfare if the trade was carried out. We incorporated the naivety factor into this checking stage and it acted as an offset. Now instead of checking the welfare after a trade, we would check the resulting welfare after the trade plus this offset. This means that agents are able to accept trades that are not necessarily to their benefit as they perceive every trade to be slightly better than it actually is. In essence, this means that good trades look better and worse trades also look better. This naivety factor which was assigned was defined as a real value in the range $[0, 0.9]$. This value as per usual would be carried by parents to their children with a 10% chance of mutation. This was done to simulate how children can learn from their parents. It was also done to determine if through evolution we could evolve to an optimal value. Below is the result of the experiment incorporating this naivety factor's evolution.




Figure 7.10: Average Intelligence for Trade Model

In Figure 7.10 we can see that the naivety factor tends to zero, zero being the point which an agent has complete understanding and knowledge of the trade in front of them. The reason it does not reach zero in Figure 7.10 is partly due to the mutation chance of the naivety factor as well as the fact that it takes longer than the simulation length to converge to zero. This value can also not become negative as a result of how this parameter has been implemented with the lower bound being zero. Allowing them to make a perfectly informed decision on whether or not to trade. The reason it converges is the same reason our metabolism converges to its low: agents who perform better (in this case agents with low value for their naivety parameter) live longer and are able to carry this trait down to their children. On the other hand agents with a higher value perform worse, die out and cannot carry their gene trait down.

By converging back to zero we essentially have the same experiment as before because now our offset parameter has little to no effect on the perception of the trade. However, this was not a pointless experiment as we have shown how we can incorporate a hidden information or naivety element into our trading rule by making a small adjustment. It is also important to note that this is not a learned strategy but rather the agents through interacting with each other and the environment evolves to be smarter.

7.5 Trade Prices

In order to set up a trade, actors (people) need to know what it is going to cost them to conduct this trade. This is apparent in all forms of trade, in most cases it is the amount of money required for goods or services. In our case we are trading one good for another, so the trade price would be the quantity of the goods you are trading in exchange for another. The price of goods and services have always been dynamic and are dependant on a number of changing elements one being scarcity. In this section, we will analyse how trading at an individual level changes over time by examining the average trade price across the world. Previously we also proposed that the laws of supply and demand were prevalent in our simple model. Therefore, now that we are analysing price, we can take a deeper look at whether this fundamental concept of economics emerges.

In order to get a better understanding on how the trading of agents evolves over the course of the simu-

lation, we decided to record all the prices p for each of the different types of trades e.g. (sugar \rightarrow spice) and averaged them across every episode. The graphs below represent these average trade prices. For example for (sugar \rightarrow spice), we read the value on the graph as the amount of sugar you are willing to trade for 1 unit of spice. This can be considered the trade price of sugar for spice.

Before we can analyse the change in trade prices we must first understand the concepts of supply and demand [Investopedia 2018a]. Demand refers to the quantity of a product or service buyers are willing to purchase at a certain price. Supply refers to the amount of goods or services producers are willing to provide at a certain price. In economics, the price is therefore considered a reflection of supply and demand. This means that instead of an individual or government which has the potential for corruption, the invisible forces of the market control pricing.




Figure 7.11: Average trade prices for Trade Model

We note initially the trade prices all start very erratic, however, they all converge to a stable level. These stable levels continue unchanged for the rest of the experiment. We note that from the way that the resources were generated in the world, there is exactly the same amount of initial sugar and spice in the world. The only difference is the locations of the different yields. The amount of food available is higher than that of both sugar and spice. Taking this into consideration, as well as the levels at which the price converges, we can generate some conclusions. Namely, we see that the price of food for spice and sugar is higher than the equivalent prices in the opposite direction. Therefore agents are more willing to give up a higher amount of food for either sugar or spice, as well as the world having the potential for more food at any time. This implies a relationship between the scarcity of the resources and the trade prices. Here we can conclude that what emerges from a simple bartering process between individual agents, is that of a global market structure whereby the prices are dependant on the demand for that resource. This emergent behaviour resembles that seen in modern society, whereby certain resources such as gold have a higher relative price based partially on scarcity. This relationship ties back to the emergence of supply and demand which was explored earlier.

In order to test to see if our experiment was, in fact, simulating a proper resource market, we decided to introduce two dynamic events. On episode 300 we reduced the capacity of sugar for each tile by 1 across the world. We also then doubled the capacity of food for each tile across the world on episode 500.




Figure 7.12: Average trade prices with dynamic events for Trade Model

As before we see an initial period of erratic pricing but the prices converge to a stable level. By episode 300 we see an immediate reaction to our simulated “sugar drought”. This dynamic event could also be said to represent some other event which would have resulted in a decrease in the supply of sugar. The amount of spice and food we are willing to trade for a single unit of sugar increases drastically. The converse is also observed whereby agents now are only willing to give up far lower amounts of sugar for food or spice. This means that our society is adapting to a dynamically changing environment and the average trade prices are depicting that change. These trade prices then converge to new levels based on this new world with a slight scarcity of sugar. The same dramatic change can be observed just after episode 500 whereby we simulate an influx of resources in the market. This could be represented as the advent of a new technology or a different supply channel, either of which would result in the increase in the supply of food. Again we see the trade prices adjust based on this event. The amount of food agents are willing to give for 1 unit of sugar or spice greatly increases as now agents have an excess of food they can get rid of. Once again the converse is true whereby agents only ever trade away very small amounts of sugar or spice for a single unit of food.

Our previous conclusion of the emergence of a dynamic market which is influenced by supply and demand remains intact. What is even more interesting is that these agents do not have access to any global average data nor are they aware of the global scarcity or abundance of resources. They are simply driven by a commonality being a need for survival and hence the global demand for resources will go up and down based on this need. They understand, “I” as an individual need a particular resource to survive, without which “I” will die, therefore I am willing to trade far greater quantities to obtain it. Hence we have shown that through individual trades that there is, in fact, an inherent price dependence on scarcity. We have also shown that trade prices will tend to some fixed point granted that the society remains relatively constant. We have also shown that the invisible forces of the market, mainly the laws of supply and demand, determined the price of resources.

7.6 Trading with preferences

In our world, acceptance of a trade is based purely on whether or not a trade is mutually beneficial. This, however, is not always the case in the real world. Many factors are considered even before the terms of a trade are brokered. One such example would be whether or not the person or entity shares your own core beliefs. For example, an animal rights activist might have second thoughts trading with a poacher. The activist in most cases would rather trade with someone with similar beliefs even if it meant they would receive a smaller net benefit than that possible with the poacher. We also see this type of preferential trading where people have loyalty to particular businesses. This section will aim to generate a society which trades based on whether or not agents share similar beliefs.

The first technique that was implemented in order to incorporate preferences into the way agents traded among each other, was by incorporating the influence rule. This was achieved by including a new parameter called trading threshold. This threshold was used to determine if two agents were allowed to trade and it worked as follows. The number of differences between both agents' cultural parameters was calculated once again by simply increasing a counter whenever the parameter differed at the same position. This number was then compared with the trade threshold parameter. If the numbers of differences were less than the trade threshold an agent was eligible to trade. If this was true for both agents they would then proceed with trading. Since there are 9 positions in the cultural parameter, the trading threshold would be defined as an Integer on the range $[0, 9]$. For example, if an agent had a low threshold value like 1 they would be heavily biased towards only trading with very similar agents. If an agent had a high value for their threshold like 7 they would, in essence, be more willing to trade with a wider variety of agents. This parameter like many of the ones previously discussed would be carried down like a gene to their children and once again there was a 10% chance for a mutation to occur.

Two experiments were done incorporating this method and the results of these are shown below. It is important to remember how the influence rule will affect the result of these experiments. Since as the world evolves under the influence rule we see changing levels of differences between people as previously demonstrated. We have shown in Figure 6.2 that the number of differences in a population tends to a low stable point. Our two experiments were designed in order to test this effect. Experiment 1 was conducted whereby the influence rule acted normally, whereby agents would interact, share ideas and eventually converge to a stable belief system. The second experiment prevented people from sharing ideas and any newborn agent would be assigned a randomly defined cultural parameter. This was to ensure we maintained a level of diversity in our population and hence maintain a higher number of differences.




Figure 7.13: Number of cultural differences grouped by experiment

In Figure 7.13 we see an initial drop off for both experiments as the population size decreases. This initial die-off has been common in all previous experiments such as in Figure 7.7. Here we see how the experiment where we allowed the influence rule to operate, resulted in fewer cultural differences for the duration. The experiment which implemented a random assignment for everyone’s cultural parameter where there was no opportunity for trade can also be seen to result in far higher levels of cultural differences for the entirety of the experiment. This result was to be expected, however, it was important in order to ensure we had established two viable experiments.




Figure 7.14: Number of trades grouped by experiment

Figure 7.14 is similar to Figure 7.13, this time depicting the number of trades for both experiments. Here we observed that the experiment which resulted in the lower number of differences led to more trades occurring. This resulted from what we observed in the previous figure. Since there are fewer cultural differences in the population, this implies there will be more opportunities to trade. Furthermore more opportunities to trade results in more actual trades occurring.




Figure 7.15: Population grouped by experiment

The final Figure 7.15 once again depicts the benefit of allowing for and encouraging more trade among agents. The experiment which resulted in more trades occurring also results in a higher population of agents emerging as the carrying capacity of the world. This demonstrates that a society which is more open to working with each other in this particular, trading together, leads to a more prosperous society for more people. The society in which its members refuse to trade based on bias against someone, leads to a less successful population. However, biases are present in everyone and represent a complex problem. Here we have shown that by making a small adaptation of a previously utilised rule we can include this complex behaviour.

Through the course of the analysis above we have demonstrated that we can simulate a society which trades based on cultural preferences. We have also shown that when the society is culturally diverse, under this system, we see fewer trades occurring. Finally, we have shown that there is an increase in the amount of trading when the society has converged to a relatively uniform belief. At this point, the experiment behaves in the same fashion as previous experiments which did not include preference. We can see that if we take a utilitarian approach, we would encourage members of the society to ignore personal preferences. By a utilitarian approach we mean that if we were to just focus on the most successful approach which results in the highest population levels. This would mean far greater utilisation of the Trade Rule which has already been shown to be beneficial. However, this forgoes all valid moral or ethical arguments which would prevent trade in some instances.

7.7 Trading based on ideology

The idea behind this section is very similar to the previous as we are still considering who agents are allowed to trade with. In this section, we consider an archetypal approach to trading which is based on groups which remain fixed throughout the lifespan of an agent. The archetype associated with an agent can be considered their personality. The goal here is to determine the differences in success rates based on your archetypal personality.

This second technique that was used in order to incorporate archetypal preferences into trading was that of trading strategies. Trading strategies was an additional parameter that was added to each agent,

along with another called trading tag. These two parameters worked in conjunction with the trading rule in order to simulate a more realistic trading environment where agents don't all behave the same. An agent's trading strategy in $\{1, 2, 3, 4\}$ determined the agent's trading behaviour as follows.

Trading strategy (Archetype)	Implication
1	trades with anyone
2	trades with like tags
3	trades with differing tags
4	doesn't trade with anyone

Table 7.1: Implications of trading strategies

An agent's trading tag was defined exactly the same as in $\{1, 2, 3, 4\}$, this value determined which group an agent was part of. Like other parameters, such as vision, both the trading strategy and tag are carried down through reproduction. A slight adjustment was made to the trading rule in order to incorporate these additional parameters. Instead of simply trading with all agents in the current agent's vision, we now first checked if both agents were compatible to trade. If they were, the trading process would continue normally as described previously. If not these two agents would not trade and the current agent would move onto the next agent. To determine if two agents were compatible, we checked each agent to see if they were able to trade with the other, based on their trading tag and trading strategy. If both agents were able to trade, they would, if even one of them could not this would result in ending the trading process. A couple of examples of how this compatibility check worked between A and B are as follows.

A Tag	A Strategy	B Tag	B Strategy	Outcome
1	1	3	3	Can Trade
2	2	1	1	Can't Trade
2	3	2	2	Can't Trade
4	4	1	2	Can't Trade

Table 7.2: Compatibility check examples between A and B

If we consider the first example on row 1 of Table 7.2. The reason this resulted in a "Can Trade" scenario was because A's strategy of 1 means that it may trade with anyone regardless of their tag. B's strategy of 3 means that both tags must differ in order for them to trade, which is true in this case. Therefore both agents can trade. If we consider the example in row 2 of Table 7.2. Here B is the same as A in the previous example so we know B can trade. However, when looking at A, it has a strategy of 2 meaning both A's tag and B's tag must be the same as A to trade. This is not the case as A's tag is 2 while B's tag is 1, therefore the outcome results in these two agents being prevented from trading. An experiment was set up in the same fashion, as previously outlined in this section, with these additions and results are shown below. The basic logic is that strategy 1 implies the agent can trade with any agent no matter the tag. Strategy 2 implies an agent can only trade with other agents of the same tag. Strategy 3 implies that an agent may only trade with agents who have a differing tag. Finally, strategy 4 implies an agent cannot trade at all no matter the tag of the other agent.




Figure 7.16: Number of agents divided by their trading strategy

In Figure 7.16 we can see the number of agents divided up into groups based on their respective trading strategies. Our initial observation was that agents with a strategy of 1, meaning they could trade with anyone, appeared in the greatest numbers. In previous experiments, we have shown that trading results in a positive increase in both wealth and numbers within a population. Therefore a strategy which resulted in increasing the number of trading opportunities would lead to this same benefit. Since being able to trade with anyone will result in the highest chance of trading with others, we concluded that these agents would receive the greatest benefit from trade and therefore appear in the greatest number. An unexpected occurrence was that of the early dominance of strategy 3 equal to that of strategy 1. We deduced that strategy 3 was performing the same because in the early stages of the simulation there is a large variety in the agents living. This means that there will be a large portion of agents nearby to those with strategy 3 who are of a differing trading tag.

Therefore, in this early period, we will see a lot more of these strategy 3 agents trading. This continues to happen until around episode 300 where the number of agents stabilises. The reason this occurs is that trading is making these agents wealthy enough to reproduce more than strategy 3 and 4. Now since these agents are reproducing their offspring have a high chance of sharing their trading tag. These offspring will also be placed either neighbouring themselves or their partner which is in turn within the current agent's vision. This leads to a world where agents with strategy 3 are increasingly going to come into contact with other agents with the same tag as themselves. This leads to a decrease in trading, and in turn, agents are less successful and reproduce less. From this process emerges a stable population where enough people of differing tags are present to support this population and there aren't too many agents with the same tag preventing agents with strategy 3 from trading and being successful. We observed agents with strategy 2 grew slowly but consistently throughout the experiment. We concluded that this was because this strategy requires agents with the same tag to be nearby in order for them to trade. If these agents exist nearby they will trade and therefore be successful. These successful agents will then reproduce and create additional agents nearby. These additional agents will have the same tag as their parents and therefore promote further trade with strategy 3 agents. This logic is similar to that of the previous observation. Since these agents require same tags for trade the more they reproduce the more they will have an opportunity to trade and so the cycle repeats. This is why constant growth is present. The reason we do not see an exponential growth is due to the limitation of the world itself. This world

doesn't have an infinite supply of resources. From previous experiments, we have seen the negative effects of not trading and that result is once again validated by agents with strategy 4.

In conclusion, the best strategy or most successful archetype which emerged from our trading rule under these trading strategy restrictions was one that allowed the agent the highest opportunity to trade. This strategy is the trade with any one strategy or number 1. Through our analysis, we have shown once again that we can generate a society which takes into consideration personal preference when it comes to trade. Finally, we have also shown that the different archetypes that were defined, result in very different success rates.

7.8 Conclusion

In this chapter, we set out to show that we could expand upon the work done by [Epstein and Axtell \[1996\]](#) by introducing a third resource. In this regard not only did we show that it resulted in similar behaviours but we also discovered by introducing an additional resource we increased the difficulty of living within the environment. We found that societies which were not able to survive the early periods had a greater chance of becoming extinct. Survival was only possible after the society had evolved to have strong performing genes. We also intended to demonstrate the benefit of introducing the ability of agents to trade on a one-to-one basis. This was shown through comparison with a no-trade strategy, whereby trading resulted in higher average age as well as population size. We did, however, observe that the average wealth was lower. We concluded that this result was due to a larger population size competing for equal levels of available resources. In both cases, the wealth is shown to be higher than the minimum required for life.

Thirdly we wanted to show that we could incorporate imperfect trading into our environment. To this end, we demonstrated that by introducing an offset to how agents perceive trades, we could, in fact, simulate imperfect information or naivety. We also showed that given a long enough period of time, the society would converge to a point where the information was in fact perfect. This meant that our population through evolution was becoming smarter. Finally, our last goal was to show that different trading strategies would result in different behaviours as well as societies which perform at different levels of success. To achieve this goal we implemented systems whereby trading was not only based on net benefit but also by individual preference. Through this process, we observed a reduction in trades and in turn the benefit of trading while cultural differences were high. Leading us to conclude that in terms of raw benefit to population sizes, trading without personal preferences would be beneficial. When we looked at archetypal trading, we also observed a similar result whereby the strategy which resulted in the highest opportunity for trading to occur performed the best. As a final conclusion, we can safely say that not only does trade allow for the existence of societies to survive in harsh environments, but it also allows them to thrive.

Chapter 8

Tax

8.1 Introduction

In previous chapters, when modelling wealth, we looked at the societies as a whole where we only considered the average wealth of the population over the course of a simulation. This approach, however, does not reveal insights at an individual level. By analysing the wealth of population at an individual level, we can then determine the distribution of that wealth among its people. This allows us to answer questions based on that distribution, such as do we have a large population of poorer people over richer people. This scenario can be seen in most modern economies which are concerned with producing goods and services for consumption. It is also generally accepted that extreme cases of this wealth inequality are bad for a society, as it means there are large portions of the population struggling to survive while a select minority are prospering. Many protests over the years have been fuelled by wealth inequality. Most notable of those was the Occupy Movement which began in New York on 17 September 2011, [Costanza-Chock 2012]. This initial protest was subsequently backed up by similar protests across 82 countries including South Africa.

From the sheer number of countries involved, we can already get a sense of how many people wealth inequality effects. It is therefore important to find fair solutions which do not inadvertently destroy the society they intend to help. To date, taxation has been used to combat this issue as one of the effects of taxation is to redistribute wealth from the rich to the poor. One could argue that the fairest distribution would be one where everyone receives an equal share, however, a counter-argument would be that the poor deserve a greater share in order to catch up. Finally, a third argument could be made that no taxation is best as those who are rich can afford to make more investments in others. Therefore, the goal of this chapter is to conduct analysis on the wealth distribution of the model generated in Chapter 7. Secondly, it will be to take that knowledge and compare it to the distributions generated by societies operating under different types of tax plans.

In the interest of implementing such a model, we must first understand what taxation is. Taxation refers to compulsory or coercive money collection by a levying authority, usually a government. Taxation is differentiated from other forms of payment, such as trade, in that taxation does not require consent and is not directly tied to any services rendered. The government compels taxation through an implicit or explicit threat of force. Taxation is legally different to that of extortion because the imposing institution is a government, not private actors. Tax systems have varied considerably across jurisdictions and time. In most modern systems, taxation occurs on both physical assets, such as property and specific events, such as a sales transaction. The formulation of tax policies is one of the most critical and contentious issues in modern politics [Investopedia 2018b].

We were able to use this knowledge in order to implement a tax based system. This required the implementation of a government object which would act in the same fashion as a government in real life although with reduced functionality. The way in which the government functions is as follows:

- At the end of each agent’s turn, the government collects and stores separately a portion of each of that agent’s resources. This is based on the value of the tax parameter. For example, if the value was 0.1 the agent would pay 10% of the amount of each resource to the government. However, if an agent has below a certain prescribed amount for a resource they do not need to pay. This is determined by the min parameter. For example, if an agent has 3 sugar and the min value is 5, they do not need to pay any sugar. This comparison is done for each resource independently and is to ensure that poor agents do not die as the result of paying tax.
- At the end of an episode (i.e. after all agents have been processed), the government will discard a portion of the total collected resources. This is to simulate the operating cost of government as well as wastage and was set at 10% for all experiments involving taxation. The exact amount a government spends on operating costs as well as the loss through wastage is difficult to determine. Hence the value of 10% was simply chosen to incorporate an extra level of realism to the model. This parameter does not include the resources that get used for example in improving infrastructure. This is because those costs are what represent the value of which the people in the society actually benefit from. Hence to simulate the benefits government provide, for example, welfare and infrastructure improvements, the government will redistribute resources directly back to the agents. One type of redistribution method will be to distribute the remaining resources equally to each individual.

There are, however, detractors to excessive taxation as they theorise that when you simply transfer wealth from the rich to the poor, you make a society less productive [Downey 2012]. This is because the rich are the ones who use their money to make investments in the economy. This makes sense in two ways, firstly if you increase taxes on the rich, you, in turn, reduce the incentive to become rich as it is less worthwhile. Why should someone work harder if it results in them paying more taxes? Secondly, the rich like everyone keep their money in the bank. The bank, therefore, has more money available to make loans. This is just one such way that the wealth of the rich impacts society. It is for this reason that we also implemented a rule which governs how agents invest back into the society. This InvestBack Rule is handled as follows:

- For each agent do:
 - Look out as far as the vision permits using Breadth First Search (same as previous rules).
 - Keep track of all agents within that vision.
 - Calculate how much needs to be invested back using Equation 8.1 for each resource.
 - Evenly distribute that amount for each resource to each agent we kept track of.

$$InvestBack = \frac{res * pop}{Tres} \tag{8.1}$$

where

res is the amount of a particular resource the current agent has.

Tres is the total amount of a particular resource the entire population has.

pop is the total population size.

The idea behind Equation 8.1 is that the richer you are, the more you will invest back into those agents around you and was inspired by the work of Downey [2012]. The experiments to follow make use of both the government structure, InvestBack Rule as well as the base model that we have been using already.

8.2 Wealth distribution control study

This chapter revolves around a deeper look into the individuals present in the society rather than previously where we looked at how the society as a whole changes. More specifically this chapter looks into the wealth distribution of agents at particular points in the simulation. This type of analysis was not done with respect to our previous experiments. It is, therefore, the goal of this section to formulate a base, by studying the wealth distributions observed in our model with and without trade in the absence of any taxation.

To study the wealth distributions we ran the same format experiment when we were comparing trade with no trade. However, now we recorded every agent's wealth across all 50 simulations at 2 points in time. These time-stamps being episode 0 and 1000. We did consider other time-stamps, however, they were relatively similar to episode 1000. It is important to note that the initial endowment was implemented in the same fashion as in Chapter 7 where agents initially received a value on the range [5, 45] for each resource. The wealth of all these agents were then binned into wealth brackets and the results are shown in the figures below.

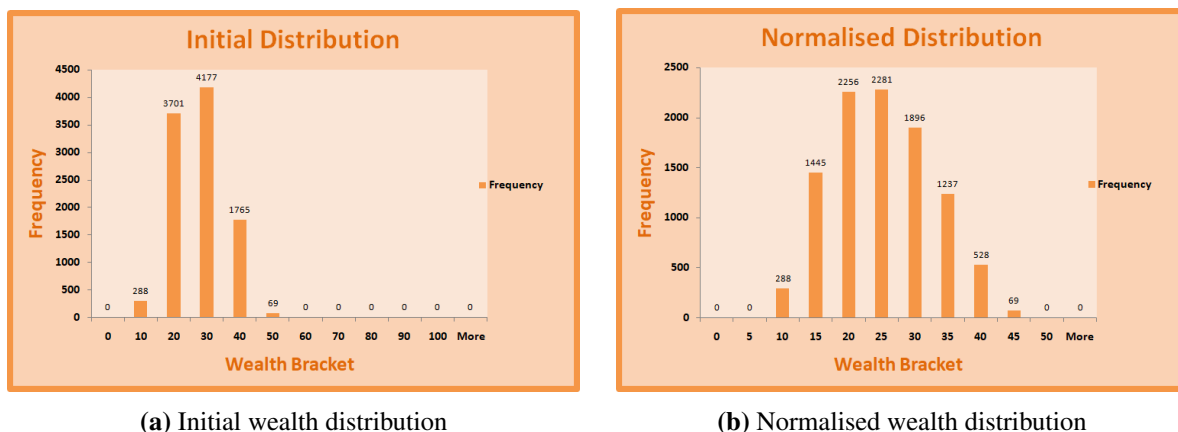




Figure 8.1: Initial wealth distributions

Firstly we looked at the wealth of our initial population of agents to get a better understanding of the initial wealth distribution. Figure 8.1a is the initial wealth distribution binned into wealth brackets of 10. This scheme was chosen as it results in a better visualisation of the wealth distribution as the population evolves. The reason this does not look like a normal distribution is due to the binning process and the range in which our values are initialised. For example wealth bracket 10 contains 6 different wealth values {5, 6, 7, 8, 9, 10} whereas bracket 50 contains 5 values {41, 42, 43, 44, 45}. This difference in possible entries is the reason why the initial distribution does not have the appearance of being a normal distribution. However, we also generated a normalised bin scheme which can be seen in Figure 8.1b. Here we binned the initial population into brackets of 5 and as can see what resulted was a normal distribution. It is important to note that the initial distribution will be the same for all experiments to follow.

The first experiment we conducted was to compare the wealth distributions between the “trade” and “no trade” experiments.



(a) Episode 1000



(b) Episode 1000

Figure 8.2: Wealth distributions of control experiments

In Figure 8.2a we see that the distribution shifts to the right. This indicates that when left to simply collect and consume resources, the population that emerges is one where there is a widespread of wealth. There are also no agents at immediate risk of starvation which can be seen by the empty first bin. Now an interesting observation was made when comparing Figure 8.2a and Figure 8.2b. That being said, we see a larger number of agents with wealth of 40 and above in the “no trade” experiment than in the “trade” experiment. More specifically preventing trade leads to approximately 29% of the population having a wealth greater than 40 compared to only 14% when trade was allowed. This, like we have previously concluded, is because there are far more agents in the trade experiment competing in the

same environment with limited resources. This means that there is far greater opportunity for particular individuals to become rich in the “no trade” experiment but this is only because the population size is so much smaller.


Not only have we shown that not all individuals live successful lives, independent of genes, but also that the Trade Rule we have implemented results in far fewer agents with wealth in the upper quartile when compared to the experiment without trade. The reason why we say that these agents are independent of their genes is by the fact that come episode 1000 all agents present have the optimal strain. In conclusion, the distributions that we observed above, represents a starting point of which is used for distinguishing the effects of taxation in the sections to follow.

8.3 Taxation versus Investment


The previous section formed the basis of our look into wealth distributions and it is this that we will build upon here. In this section, we compare our taxation strategy as well as the InvestBack Rule with trade in order to determine the strategy (or combination of strategies) which benefits the population the most. In the figures below we utilise particular keywords to indicate different conditions for experiments. The meaning of each of those is as follows:

- Tax: This indicates that taxation was present at 10% of an agent’s wealth where the government wasted 10% of its total collection and equal redistribution occurs.
- Invest: This indicates an agent would invest in the neighbouring agents using the InvestBack Rule.
- Trade: This indicates that agents were allowed to trade using the global trade rule.


The figures representing the wealth distributions of our experiments using the different combination of strategies can be seen below.




(a) Episode 1000



(b) Episode 1000



(c) Episode 1000



(d) Episode 1000

Figure 8.3: Wealth distributions of differing tax strategies

The first interesting observation to be made is when taxation is involved as seen in Figures 8.3a and 8.3c. Here we see that roughly 99.5% of the population is in bins 10, 20 and 30. Therefore due to the only difference between the two experiments being the ability to trade, we can conclude that trade promoted the increased wealth. The opposite effect can be seen in Figures 8.3b and 8.3d where no tax is present. Here we see agents are able to achieve far greater wealth than that possible in the tax experiments. We also see that there are far fewer agents in the second smallest bin. This leads us to conclude that not only does tax reduce the ability to earn greater wealth but it also leads to more poorer people. When comparing Figures 8.3b and 8.3d, we see that when agents only invest back there are fewer poor people. However, there are also fewer people with wealth 40 and above.

Now in order to get the complete picture, we need to understand the total population size represented by these distributions. This is because a rich society of 10 people is not necessarily as successful as a poorer society of 100 people. The figure below represents the total population over time for each experiment discussed above.




Figure 8.4: Population for different wealth redistribution strategies

From Figure 8.4 we can see, although our previous figures would lead us to believe that taxation was a detriment, the largest population emerged from the “Tax + Invest” strategy. We also observe that the addition of investing back into the population increases the population when it is in conjunction with trade. When investment was the sole form of wealth distribution it performed as poorly as when no trade was allowed. It is important to note that the “no trade” experiment also did not include investment or taxation. A final observation was that all the strategies which included trade tend to result in very similar levels of population. This leads us to conclude that the inclusion of trade was the biggest driver for the level of population.

In conclusion, we can state that although taxation does indeed lead to a higher population it does also drastically limit the opportunity for individuals to become wealthy. The antithesis of this result was also seen when taxation was not present in the experiment. Here we noted that without tax, individuals had a freedom to become wealthy. We also saw that the average wealth in these cases was also higher than that when taxation was present. These results seem to back up the theory proposed by the Occupy Movement detractors. This is especially apparent when we compared “Only Invest” with “Tax + Invest”. Here the “Only Invest” strategy even resulted in fewer poor people and much more people with a wealth in the higher bins. Whereas nobody had a wealth greater than 40 under the “Tax + Invest” strategy. The limiting factor in our experiment is the fact that wealth generation is directly dependant on the environment and the rate at which an agent can collect resources. This means that in an environment where wealth creation is less constrained, we could see a greater possibility for wealth generation. This is important because these more wealthy people would be able to re-invest more back into the economy or environment.

8.4 Different Tax rates

In the previous section, we were comparing how our taxation strategy compared to that of investment and trade. In this section, we take a deeper look directly at our method of taxation. Now we look at running

experiments with different values for the tax rate in order to determine at which levels taxation performs best. It is important for us to determine optimum tax rates which both promotes larger populations with greater individual wealth. This combination of factors represents the most successful society. Finally, we must not forget that in the previous section taxation was only experimented with in conjunction with the InvestBack Rule. In this section, we isolate taxation for further analysis in order to determine if this combination of tax and investment is the reason for the poor performance.

In order to accomplish this goal determining optimal tax rates, we ran separate experiments with differing tax rates. These being 5%, 10%, 20% and 30%. Our wastage or government operating cost parameter was set at 10%. The exact amount of money a government loses is very difficult to calculate and as a result a reasonable value was chosen. This percentage doesn't change across any experiment and was incorporated to show that, one can include this parameter and yet still obtain a stable society. Finally, we disabled the InvestBack Rule as well as the Trade Rule. Therefore, the figures below are the results of these experiments where only the tax strategy is impacting the population




Figure 8.5: Wealth distribution for different tax rates

In contradiction to the previous results obtained with regard to taxation strategies, here we observe that taxation alone doesn't prevent agents from becoming rich. Rich agents being those who are able to obtain a wealth of 40 or higher. When comparing the number of rich agents for each of the different tax rates, we see that the lower tax rate leads to higher numbers of rich people. By taking our definition of rich for Figure 8.5a we get 1555 people and in Figure 8.5d we get 1004. We did note that in Figure 8.5c there were only 967 rich people, however, we concluded that this small difference was due to simply

taking a single point in time. We also hypothesise that tax rates higher than 20% impact the society less and less. This is due to the fact that the majority of resources collected by the government does find its way back to the agents. If we were to increase the operating cost parameter, we would see a far greater impact from higher tax rates. When comparing the other end of the spectrum we see relatively the same amount of agents with a wealth of 10 or less.

Once again, to get a true sense of the success of a tax rate we need to look at the population size that emerges in the presence of these different percentages. Figure 8.6 below represents the total population for each of the tax percentages discussed above.




Figure 8.6: Population for different tax percentages

Here we see that the society under a 5 or 10% tax rate converges to a higher level than that of 20 or 30% which performs equally.

Now taking that in conjunction with our previous observation of all tax rates have relatively equal amounts of poorer agents, we can conclude that higher tax rates negatively affected the society. We can also conclude that lower tax rates are more beneficial to society since these rates result in richer agents on average and larger populations. We have therefore shown that the optimal tax rate in our experiments is that of 5%. We also set out to study taxation in isolation. To this end, we can conclude that not only does taxation lead to poorer societies but also smaller ones. This can be seen when comparing Figure 8.4 in the previous section to 8.6 shown above. Here we can see taxation in isolation performs categorically worse than the strategies where trade and investment are involved.

8.5 Progressive Tax Rates

The previous experiments have been conducted using a system where all agents pay a uniform percentage of their wealth. This means that richer agents still pay more in tax than poorer agents, however, this is not how taxation occurs in the modern world. What we actually see are progressive tax systems where people are divided up into tax brackets based on their wealth. These different brackets have different tax rates associated with them. This means that richer people will fall into the higher brackets and hence

will be charged a higher percentage of their wealth. This percentage gets progressively smaller the lower the tax bracket a person falls under. The goal of this section is to implement different progressive tax systems and analyse the effect on the wealth distributions of the society.

In order to implement this new system, we needed to first define different tax brackets as well as a way of determining the bracket to which an agent belonged. For this, we decided to place agents into brackets based upon how many standard deviations away from the mean wealth they were. We then defined tax rates for each of these brackets. These brackets were defined in a way in which there was increasing pressure placed upon the richer agents. It is important to note that for each agent, we handle the wealth of each resource independently. This means that an agent with a high amount of sugar and low food will be placed into a high sugar tax bracket while also being in a low food tax bracket. It is also important to note that the standard deviation and average wealth, for each resource, was calculated each episode in accordance with the changing population. Like previous experiments agents will not pay tax if it means risking their lives. Table 8.1 below depicts the three different tax plans that were defined, where σ represents the standard deviation, μ is the mean wealth and x is the agent's wealth.

Plan	$x \leq \mu - 2\sigma$	$\mu - 2\sigma < x \leq \mu - \sigma$	$\mu - \sigma < x \leq \mu$	$\mu < x \leq \mu + \sigma$	$\mu + \sigma < x \leq \mu + 2\sigma$	$\mu + 2\sigma < x$
1	2%	5%	10%	15%	20%	25%
2	2%	5%	10%	20%	30%	40%
3	2%	5%	10%	20%	40%	60%

Table 8.1: Progressive tax plans

We ran individual experiments using each of the three tax plans above. Our wastage parameter was set at 10% with the minimum wealth required to pay tax being 5. Our new taxation system works as follows:

- The initial average wealth and standard deviation for each resource are calculated during initialisation.
- During each agent's turn, they are placed into a bracket for each resource based on 8.1. They then pay a percentage of their wealth to the government object based on their bracket.
- Once each agent's turn has concluded, the government redistributes the wealth equally back to all members of the population.
- Finally the average wealth and standard deviation for each resource is recalculated

After conducting an experiment with Tax Plan 1 we observe that there is a very low average population as seen in Figure 8.7 below. After examining the visualisation we witness that many of simulations had populations which became extinct. There are two reasons for this, one is that the agents with the worse genetic traits are not dying out quick enough. The second is that the agents with the better genetic traits are not gaining enough wealth to reproduce effectively because they are paying too much in taxes. We hypothesise that this new tax plan system is applying too much pressure to the population during the early stages when it is still evolving. We do, however, conclude that this tax system negatively affects the early stages of development in a society. This can be seen when comparing the population sizes from Figures 8.7 and 8.4.




Figure 8.7: Population for tax plan 1

In order to prove our hypothesis we ran the same experiment, however, we only introduced our taxation system after episode 200. This means that before that point agents are only collecting and consuming resources in the same fashion as the “No Tax + No Trade + No Investment” experiment. This was conducted in order to negate the effect of taxation during the crucial early stages where agents are evolving to better survive in the environment. We analyse the Population and Total Births in Figures 8.8 and 8.9 respectively to get a better understanding of the viability of the societies under the different tax plans.




Figure 8.8: Population for tax plans where tax was introduced after episode 200

In Figure 8.8 we first observe the initial die off and slow regrowth period. This continues until episode 200 after which we see an immediate population boom. This population explosion is a result of the

wealth being redistributed to poorer agents. These poorer agents, in turn, become wealthier and therefore there is a higher initial chance of reproduction being possible. Previously, we have demonstrated the disadvantages of population explosions as it leads to a population which ages at the same time. We observe the result of this ageing population once again in Figure 8.8 where around episode 290 a population collapse occurs. All three tax plans follow the same trend up until this point after which they converge to very different stable levels. Tax Plan 1 results in the highest level around 70. This is the only plan that can be considered viable as Tax Plan 2 and 3 results in a population which is far lower than any scheme seen in Figure 8.4. We also observe that many of the 50 simulations which are averaged to form an experiment result in the population completely dying out.

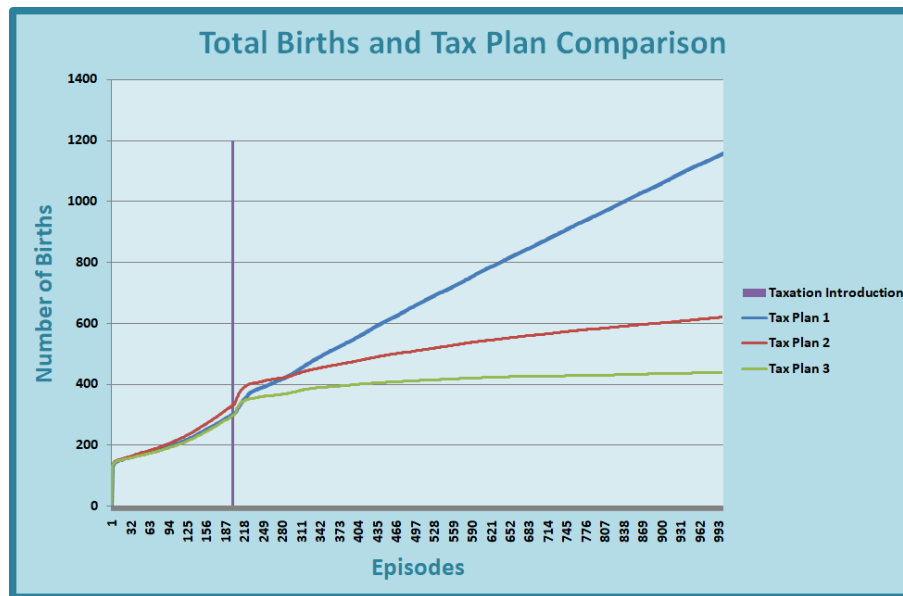


Figure 8.9: Total Births for tax plans where tax was introduced after episode 200

In Figure 8.9 we consider the total cumulative births recorded at the beginning of each episode. Here we observe how Tax Plan 1 has the greatest birth rate followed by Tax Plan 2 and lastly Tax Plan 3. We also note that after the advent of taxation we observe a dramatic increase in the number of births after which the rate stabilises. We recognise that Tax Plan 2 and 3 have relatively small gradients. This means that very few new agents are being created over the course of many episodes. This is unhealthy for a population as it could result in extinction events when the majority of the population becomes too old to have new children. The extreme case of this is seen in Tax Plan 3 where the gradient is almost 0 meaning no new agents are being born. This trend continues for many episodes and signifies how unviable this plan is. We also note that Tax Plan 1 which has the lowest tax rates for richer agents results in the most successful society. Furthermore, as the progressive tax rates become larger for Tax Plan 2 and 3 so does the success of the society drop.

We hypothesised that these tax plans were applying too much pressure on the population during the early stages of the experiment. Here we have shown that is not the only reason for the failing of the tax plans. We do note that delaying the period in which taxation was introduced, helps the success of a society to some degree.

Based on these results we looked to determine if there was a different way to increase the viability of these progressive tax systems. Therefore, we decided to change the cost incurred to parents for having children. Previously when determining whether reproduction would occur both agents were required to contribute half of their own initial resource value for each of the respective resources. This parameter

was called the initial endowment. This was considered the cost of raising a child to be a fully functional member of society. In the following experiment, we tested each tax plan where the endowment (cost onto each parent) would be fixed at 5. This was done in order to promote additional reproduction for the purpose of increasing the chances of viability when using our progressive tax system. The value of 5 was chosen as it is the lowest possible value an agent can be born with based upon our initial distribution of resources being on the range [5,45]. Therefore was chosen a best case scenario when it comes to the cost of reproduction. It is important to note that this experiment did not delay the introduction of taxation to episode 200 but rather had taxation active from the onset.




Figure 8.10: Population for tax plans with low initial endowment

We observe in Figure 8.10 that by decreasing the initial endowment required we were able to increase the population for Tax Plan 1 and 2, however, Tax Plan3 still resulted in an unviable society. For Tax Plan 1 we see the population that emerges is almost double that we observe in 8.8. However, the benefit was not as great in Tax Plan 2 where the emergent population level is only slightly greater than that in 8.8. We also decided to consider the wealth distributions which can be seen below.



(a) Episode 1000



(b) Episode 1000



(c) Episode 1000

Figure 8.11: Wealth distributions of Tax plans with low initial endowment

Firstly we recognise that the population size is too small under Tax Plan 3 in Figure 8.11c for it to be comparable to the others. However, when comparing Tax Plan 1 and 2 in Figures 8.11a and 8.11b we observe that Tax Plan 1 results in the majority of the population being in a higher bin to that of Tax Plan 2. Couple that result with the observation that Tax Plan 1 leads to a higher stable population we can conclude that Tax Plan 1 was the most beneficial to the society. This is interesting as it is also the plan whereby the rich agents are taxed the least.

Once again we observe that not all of the tax plans resulted in a viable society. It is for this reason that we conducted another experiment which utilised both the changes in the previous two experiments. Those

being the delayed taxation to episode 200 as well as the reduced initial endowment (cost of reproduction for parents).

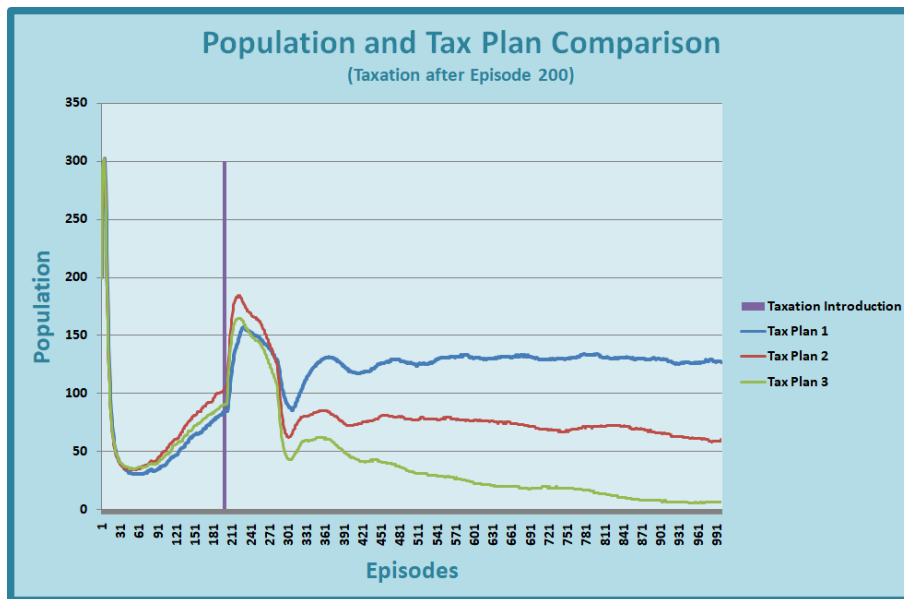


Figure 8.12: Population for tax plans with low initial endowment and taxation only after episode 200

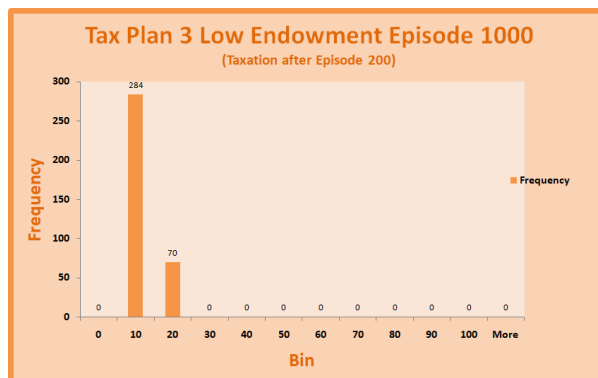
In Figure 8.12 we observe many of the same phenomena as seen in previous experiments, most notably is that of the population boom after taxation is introduced and the subsequent die-off period. We also observe that the combination of both factors (delayed taxation and low endowment) have resulted in a cumulative benefit for all tax plans studied. All tax plans lead to a population which survives to the end of the experiment. However, given a longer period of time, we will most likely see a further decrease in the population sizes for both Tax Plan 2 and 3. Tax Plan 1, however, remains stable around 130.



(a) Episode 1000



(b) Episode 1000



(c) Episode 1000

Figure 8.13: Wealth distributions of Tax plans with low initial endowment and taxation only after episode 200

Once again we compare the wealth distributions in order to better understand the quality of life. We observe in Figure 8.13a that 63% of the population is in the higher wealth bin. A very similar observation can be made for Tax plan 2 in Figure 8.13b where 61% is in the higher bin. If not for the difference in overall population size one might consider these plans equivalent. Now when considering Tax Plan 3, the opposite can be seen in Figure 8.13c where 79% of the population is found in the lower bin. A final observation is that for all experiments utilising our progressive system, none of the wealth distributions observed had agents with a wealth value greater than 30.

Now an important part of our system to realise is that as the average wealth of the population shifts so

does our tax rate system. This means that even if you have a really low average wealth you will still find agents in the upper brackets of taxation. This means that an agent who would be considered objectively poor is still charged at a high tax rate due to there being many poorer agents.

Therefore we conducted a final experiment utilising a new progressive tax system where the brackets were fixed. Once again in Table 8.2 below x represents the wealth of an agent. Now regardless of the wealth of the agents, the tax brackets are fixed. Therefore we will no longer observe the scenario previously described where poor agents are paying high tax rates. Once again each resource is handled independently, therefore, agents might be in a high tax bracket for one resource while being in a low tax bracket for another.

Plan	$x < 10$	$10 \leq x < 20$	$20 \leq x < 30$	$30 \leq x < 40$	$40 \leq x < 50$	$x \geq 50$
1	2%	5%	10%	15%	20%	25%
2	2%	5%	10%	20%	30%	40%
3	2%	5%	10%	20%	40%	60%

Table 8.2: Progressive fixed tax plans

These 3 plans were designed so that they would get progressively higher for the richer agents which contribute the most to the total wealth that gets distributed back. This meant that we could get a sense of how much strain we could put on the richer agents before it resulted in a poorer performing strategy. This scheme of the richer you are the higher you get taxed is echoed in modern society and although our percentages don't match exactly to that of a particular scheme, it follows the same logic. [SARS 2018]

We ran an experiment for each of the three tax plans described in Table 8.2 where taxation is introduced immediately to the world and the initial endowment was reset back to original state of being half an agent's initial wealth. This means that if an agent was born with 20 sugar their initial endowment for when they reproduce will be considered to have a value of 10. The Figures below are the results of these experiments.






Figure 8.14: Population for tax plans with fixed tax brackets

In Figure 8.14 we can see that the populations that emerge all converge between 150 and 200. These are


all far better than anything achieved by our previous non-fixed progressive tax system. We also observe that Plan 1 is the worst performing, which is interesting as this is the plan with the lowest tax rates on higher brackets. Now if we compare this system with our flat tax results in Figure 8.6, we observe the fixed bracket system has far superior population sizes. If we go back even further and compare the results in Figure 8.4 we see that the only society whose population grew larger was that the “Tax + No Trade + Invest” strategy. Now although that was a flat tax strategy it also included investment. We note that none of these progressive tax systems were tested with investment active.



(a) Episode 1000



(b) Episode 1000



(c) Episode 1000

Figure 8.15: Wealth distributions of Tax plans with fixed tax brackets

From Figures 8.15a, 8.15b and 8.15c we observe that each of the three plans results in similar wealth


distributions, where the majority of the population is in bin 20 with relatively equal population levels in the adjacent bins. We note that the entire population falls within these 3 bins. Once again, however, we conclude that taxation has limited the opportunity of generating large amounts of wealth.

In this section, we set out to analyse the effect of a progressive tax system where agents were taxed differently based on which wealth bracket they fell into. From the different experiments conducted we conclude that our initial approach to taxation negatively affects the success of the population. This approach was where our brackets shifted with respect to the population. Here we conclude that this shifting approach results in forcing populations to become poorer and poorer. This is due to observing far lower stable population levels than in previous experiments as well as greatly reduced opportunity for agents to get out of the lower wealth bins. This poor performance even persisted after making life easier for the agents by delaying taxation as well as decreasing the cost of reproduction. Tax Plan 1 which features the lowest tax rates resulted in being the best performing system. Our second implementation, however, produced promising results and even outperformed our fixed tax systems. This approach was where we fixed the tax brackets. Here, however, we still observe a barrier to obtaining high amounts of wealth due to taxation. Interestingly with this system it was Tax Plan 3, (the plan with the highest tax rates) which performed the best.


8.6 Indirect Wealth Redistribution

Our taxation implementation works by directly distributing an amount of each resource to all of the agents. Although this is indicative of welfare programs such as social grants for the poor or old, it does not accurately represent the most prominent forms by which government redistributes wealth back to its population. These being public services such as a police force and infrastructure development, such as roadworks. Both of these are basic examples of how citizens reap benefits from their government. Taking this into consideration, our model becomes slightly unrealistic. Therefore this section will analyse the effect of a new system whereby instead of wealth being directly distributed to the people. It will instead be fed back into increasing the resources available to collect forming our indirect wealth distribution cycle. The goal of this section is therefore to analyse the effect of such a system and whether or not it is comparable to the other strategies discussed in this section.

To implement this new system we decided not to distribute any wealth collected by the government back to the agents. Instead at episode 333 and 666 we increase the capacity of every tile by 1 for each resource. This is to simulate the idea that after the government has collected enough revenue they can implement a global policy which benefits the entire population. For example, if the capacity at a particular coordinate for sugar, spice and food was 1, 2 and 3 respectively, it would be changed to 2, 3 and 4 at episode 333. It is also important to note that we utilised the uniform taxation method at 10% rather than the progressive system. The figures below represent the results obtained from this experiment



(a) Population under environment upliftment



(b) Average wealth under environment upliftment

Figure 8.16: Population and wealth distribution for Indirect Wealth Redistribution experiment

In Figure 8.16a we note that there is a spike in population growth right after the environment is uplifted by the government intervention. After the spikes, the population converges to a stable level representing the carrying capacity of the current society. This stable level then shifts after the second government intervention. It is important to remember that this government intervention simulates infrastructure improvements such as building a new highway which makes it cheaper to travel to work. When analysing Figure 8.16b we note that the average wealth for each the 3 resources dips slightly after each point of intervention. One would assume this would go further down with the explosion in the population at these points. However, this is countered by the fact that the environment is richer. This means that we can sustain a much larger number of agents at slightly lower wealth levels.




Figure 8.17: Wealth distribution at Episode 1000 under environment upliftment

In Figure 8.17 we take a look at the wealth distribution under this upliftment system. Here we observe a similar wide distribution of wealth as that seen previously. The agents have the opportunity to gain large amounts of wealth. However, the average wealth of the society is still performing worse than the trade and investment strategies.

To conclude we have shown that it is possible to implement a more realistic indirect method of wealth redistribution. The most interesting observation that was made was that we were able to generate a society which converged to a much larger level than that of any previous experiments. However, there is an argument to be made that the government intervention we simulated results in a permanent positive effect on the society and that this permanence is unrealistic. A possible future experiment to address this would be to slowly depreciate the capacity of the resources across the environment. We have also shown that we can generate a wealth distribution similar to that previously discussed using the direct taxation method. This demonstrates that this methodology is comparable to that of the direct taxation method.

8.7 Conclusion

Taxation and the levels at which taxation should occur has also been a contentious topic as it affects all of us personally. In this section, we took a deeper look into the wealth of individuals, by first analysing the wealth distributions of our previous experiments. In this pursuit, we were able to observe the distribution of our initial population as well as how it evolved to form a wide distribution of wealth. We also set out to analyse the effect of taxation on the population as well as its effect in conjunction with other elements such as trade. We also took a look at the theory proposed that it is actually the rich who through investment stimulate and uplift an economy. To this end, we showed that there was some credibility to this theory as our experiments showed that investment resulted in fewer poor people and did indeed promote wealth generation. It did, however, lead to smaller overall population size when compared to taxation. Therefore the argument is whether a larger poorer population is better than a smaller richer population and to what extent. Furthermore, we were able to analyse the optimality of differing tax rates leading us to conclude that the lower tax rates performed better when they were fixed. Secondly, during these tax rates experiments, we were able to deduce that taxation in isolation performed poorly in comparison with other strategies. We were also able to implement two progressive tax systems one of which (fixed brackets) outperformed our previous strategies in terms of the population. However, we still observed the drawback of taxation, that being the limitation on large amounts of wealth generation. Finally, we were able to generate a system of indirect wealth distribution by simulating the transfer of wealth from government to citizens through the benefits of infrastructure development. This system was also comparable to the previous direct system. We note that this chapter does not consider tax avoidance or the psychological influences of tax. For example, the negative effect taxation has on enticing people to earn more. In conclusion, we have demonstrated that yes taxation has its role to play in a society but the detractors of excessive taxation have a point. We have demonstrated a reduction in the number of poor can result from investment as well as trade.

Chapter 9

Conclusion

The work done by [Epstein and Axtell \[1996\]](#) was revolutionary by introducing the world to the idea of modelling complex social behaviour using multi-agent systems. The research that has followed and pushed development even further forward in this field based on their initial work is proof of this. This research attempted to take this idea once again a step further, by applying this approach of evolutionary modelling of societies from the ground up. By the end of the research, we hoped to be able to deduce exact cause and effect relationships between all the experiments that have been previously mentioned as well as demonstrate the emergence of pre-existing real-world phenomenon within the society that was created.

From the results, we obtained and through the discussion and analysis of the societies evolved we believe to have achieved these aims. Throughout the discussion of our results, we have given detailed quantitative analyses of the data we have presented. Although this empirical study produces interesting results, it is not without its limitations, those being we cannot analyse the exact evolution of an individual agent. Secondly, it can become hard when we increase the complexity of rules to determine reasoning for complex behaviour. Therefore pursuing formal analyses such as theorems and proofs where possible is still equally important, however, we have shown that there are still insights to be gained from this generative modelling technique.

Firstly we were able to demonstrate that the behaviours found in the SUGARSCAPE could be replicated. From this base, we were then able to expand upon many of the ideas. These led to the discovery of interesting observations such as the emergence of Darwinian evolution. From these experiments, we saw the emergence of a dominant strain of parameters representing genes. Our model was also shown to be comparable to other traditional models like the Predator Prey model. This technique was also shown to be useful when trying to predict the effect of certain policies when we studied birth control policies.

The usage of the cultural rule formed a large component of this research. From its usage, we were able to observe the emergence of uniform belief systems along with the effect individuals had on the population. From here we expanded the idea to study the propagation of information and how it is affected by the topology of an environment. Therefore we were able to show that a large number of problems could be modelled with only slight modifications to a couple simple core rules.

Arguably the most interesting results involved trade in which we were able to generate a trade environment where the trading of agents, at an individual level, resulted in a market price forming. This market price was also affected by the supply of a resource within an environment which was seen when we artificially induced a natural disaster resulting in a decrease in supply. This relationship demonstrated the workings of the laws of supply and demand. This led us to conclude that market forces had indeed

emerged in our society as a result of our simple trade rule.

Finally, we were able to compare the effect of taxation versus other policies such as investment and trade. Here we observed results more in line with the detractors of taxation, by seeing how taxation resulted in larger poorer populations. We also analysed the effect of different tax rates, where we concluded that lower tax rates were more beneficial. Lastly, we demonstrated the benefits of using a progressive fixed bracket taxation system.

There is still a lot more potential for this technique and many more questions to be answered. Some of these being, what is the effect of harsher environments or which genes are more important than others or even what would happen if we expanded reproduction to include attraction. Further questions about trade could also be considered including introducing a black market or even expanding the intelligence factor so agents could learn to take advantage of others. All of these examples serve to demonstrate the possibilities within this existing framework called an Artificial Society.

Despite its simplicity, the model generated outcomes that resembled complex real-world phenomena. It is important to remember that in none of these experiments was the global behaviour known prior but rather evolved through local interactions of agents. Additionally all results and observations of the emergent behaviour made were from the environments and systems generated in our model and may not directly translate to the behaviours of the real world. It is, however, the similarities that can be found in our model, which can be used to further our understanding of the complex behaviours observed in the real world. In conclusion, we hope to have demonstrated the strengths of Artificial Societies by showing many different applications of this technique.

References

- [Andy George 2018] Brian Stemmler Andy George, Chris Peck. *How to Make Everything : Sandwich*. <https://htme.squarespace.com/>, 2018. Online; accessed 2018-01-17.
- [Buzing *et al.* 2005] PC Buzing, AE Eiben, and Martijn C Schut. Emerging communication and cooperation in evolving agent societies. *Journal of Artificial Societies and Social Simulation*, 8(1), 2005.
- [Cerioli 1999] Andrea Cerioli. Measuring the influence of individual observations and variables in cluster analysis. In *Classification and Data Analysis*, pages 3–10. Springer, 1999.
- [Chen 2003] XUEFENG Chen. The social impact of China’s one-child policy. *Harvard Asia Pacific Review*, 7(1):74–76, 2003.
- [Costanza-Chock 2012] Sasha Costanza-Chock. Mic check! media cultures and the occupy movement. *Social movement studies*, 11(3-4):375–385, 2012.
- [Davies and Quinlivan 2006] Antony Davies and Gary Quinlivan. A panel data analysis of the impact of trade on human development. *The Journal of Socio-Economics*, 35(5):868–876, 2006.
- [Depew *et al.* 1995] David J Depew, Bruce H Weber, and Paul E Griffiths. *Darwinism evolving: Systems dynamics and the genealogy of natural selection*. JSTOR, 1995.
- [Downey 2012] Allen B Downey. *Think complexity: complexity science and computational modeling*. ” O’Reilly Media, Inc.”, 2012.
- [Epstein and Axtell 1996] Joshua M Epstein and Robert Axtell. *Growing artificial societies: social science from the bottom up*. Brookings Institution Press, 1996.
- [Factbook 2015] CIA Factbook. The world factbook; 2010. *See also: <http://www.cia.gov/library/publications/the-world-factbook>, accessed January, 30, 2015.*
- [Gardner 1970] Martin Gardner. Mathematical games: The fantastic combinations of John Conway’s new solitaire game life. *Scientific American*, 223(4):120–123, 1970.
- [Gautam 2005] Rajesh K Gautam. The primitive tribes of Andaman and Nicobar Islands. *Primitive tribes in contemporary India: Concept, ethnography and demography*, 1:35–73, 2005.
- [Goh 2011] Esther Goh. *China’s one-child policy and multiple caregiving: raising little suns in Xiamen*. Routledge, 2011.
- [Goldsmith 1974] Edward Goldsmith. *Blueprint for survival*. Signet, 1974.
- [Hammond and Axelrod 2006] Ross A Hammond and Robert Axelrod. The evolution of ethnocentrism. *Journal of Conflict Resolution*, 50(6):926–936, 2006.

- [Hartshorn *et al.* 2013] Max Hartshorn, Artem Kaznatcheev, and Thomas Shultz. The evolutionary dominance of ethnocentric cooperation. *Journal of Artificial Societies and Social Simulation*, 16(3):7, 2013.
- [Hvistendahl 2010] Mara Hvistendahl. *Of population projections and projectiles*, 2010.
- [Investopedia 2018a] Investopedia. *Economics Basics*. <https://www.investopedia.com/university/economics/economics3.asp>, 2018. Online; accessed 2018-01-17.
- [Investopedia 2018b] Investopedia. *Taxation*. <https://www.investopedia.com/terms/t/taxation.asp>, 2018. Online; accessed 2018-01-17.
- [Klöck 2007] Daniel Klöck. *Extended Sugarscape model in XL*. Technical report, Technical report, Department of graphic systems polytechnic university of Cottbus. Online: <http://danielkloeck.wdfiles.com/local-files/extendedsugarscape/StudienarbeitSugarscape.pdf>, 2007.
- [Martin 2013] Edwin Martin. *Game of Life Java applet*. <https://bitstorm.org/gameoflife/>, 2013. Online; accessed 2017-06-12.
- [Meadows *et al.* 1972] Donella H Meadows, Dennis L Meadows, Jorgen Randers, and William W Behrens. The limits to growth. *New York*, 102:27, 1972.
- [Oliphant 1996] Michael Oliphant. The dilemma of Saussurean communication. *BioSystems*, 37(1-2):31–38, 1996.
- [Oliphant 1997] Michael Oliphant. *Formal approaches to innate and learned communication: Laying the foundation for language*. PhD thesis, University of California, San Diego, 1997.
- [Penrose 1959] Lionel S Penrose. Self-reproducing machines. *Scientific American*, 200(6):105–117, 1959.
- [Reynolds 1987] Craig W Reynolds. Flocks, herds and schools: A distributed behavioral model. *ACM SIGGRAPH computer graphics*, 21(4):25–34, 1987.
- [Rouchier 2002] Juliette Rouchier. Multi-agent systems for the study of social complex systems (a way to represent social complexity—examples and validation). In *Agent Oriented Software Engineering*, volume 1957, pages 235–251. LNCS, 2002.
- [SARS 2018] SARS. *Rates of Tax for Individuals*. <http://www.sars.gov.za/Tax-Rates/Income-Tax/Pages/Rates%20of%20Tax%20for%20Individuals.aspx>, 2018. Online; accessed 2018-07-01.
- [Squazzoni 2018] Flaminio Squazzoni. *Journal of Artificial Societies and Social Simulation*. <http://jasss.soc.surrey.ac.uk/JASSS.html>, 2018. Online; accessed 2018-07-01.
- [Von Neumann and Oxtoby 1988] John Von Neumann and John C Oxtoby. John Von Neumann. *American Mathematical Soc*, 1988.
- [Zill 2016] Dennis G Zill. *Differential equations with boundary-value problems*. Nelson Education, 2016.