

**DONUT WORLD AND THE DUOPOLY
ARCHIPELAGO: SOCIAL LEARNING AND THE
EVOLUTION OF COMPETITION.**

Huw Dixon
Chapter 8 of *Surfing Economics*.

References and bibliography for this article can be found at

<http://huwdixon.org/SurfingEconomics/references1.pdf>

<http://huwdixon.org/SurfingEconomics/index.html>

Chapter 8.

Donut world and the Duopoly Archipelago: social learning and the evolution of competition.

8.1. Introduction.

The traditional approach to economics has been to assume that agents are rational and use all of the information they have in an optimal manner. However, as we have seen in *Thoughts on Economic Theory and Artificial Intelligence* there are many arguments against this. At best, it is a modelling simplification, an “as if” assumption made to make the process of understanding the economy and economic behaviour easier.

There are *rationality fundamentalists* around, who believe optimising rationality is an essential part of human nature. I think that this is largely a credo with little or no justification, an act of faith by economists who want to have a single principle with which to understand economic phenomena. There are reasons why I reject the fundamentalists view. First, most economic decisions are made in the context of groups of people: the family/household, the firm, the union, the bank and so on. Even if individuals are “rational”, that does not imply that the decisions of groups will be “as if” made by a single rational individual. Second, in practice individuals do not appear to act in ways consistent with rationality all the time¹: they may learn to be rational, particularly in repeated situations where there is a lot to be gained or lost. But then again, some people end up making the same mistakes over and over again.

In recent years there has been a considerable revival of interest in the notions of learning in a *boundedly rational* context. This idea has of course been around for as long as economics itself. However, rather paradoxically, in the last decade the idea of boundedly rational processes has been revived in the field of game theory². Game theory has traditionally been the area of economics where the belief in rationality has been the most intense. Indeed, many game theorists inhabit an artificial world where disembodied rational agents interact in a sea of common knowledge, able to perform all and any calculation the theorist might conceive. Without any constraints on the imagination, unencumbered by notions of firms or markets or any explicitly economic context, with a fascination for 2x2 games (the Prisoner’s Dilemma an obsession) they create a *rationality wonderland*. In rationality wonderland, agents ar

perfectly rational, agents know the structure of the game and also know that all agents including themselves know that they know the game. This is called *common knowledge*. In order for a rational player to know what to do, she³ needs do two things: first to guess what the other guy(s) are going to do; second to choose a best response to that action. Now, this problem involves an infinite regress: I need to predict the other player's behaviour to choose my best action; she needs to predict mine to choose her best action. So, I need to predict her prediction of my action; she needs to predict my prediction of her prediction of my prediction etc. This infinite regress is *solved* by game theorists uttering the incantation "common knowledge" and then proposing that rational players would do what the game theorist wants them to. In my opinion the concept of common knowledge is incoherent, and arises because economists (in this case game theorists) try to extract too much from the basic idea of rationality⁴. However, dear reader, this is not the time nor place to explore this line of reasoning. Rationality wonderland is not our destination now: our destinations are *Donut World* and the *Duopoly Archipelago*. Before we set off, we will briefly consider *learning*.

Learning can be seen as taking place at two levels. *Individual* learning occurs when a single agent alters its beliefs and/or behaviour. Learning in this sense can take place if there is only one agent on its own without any interaction with another: for example, Robinson Crusoe was able to update his beliefs about farming and fishing techniques during his stay on the island. In the previous chapter on *Artificial intelligence and economics* I discussed some aspects of modelling the bounded rationality of individual agents. *Social* learning occurs within a population of agents and can only occur when there is more than one agent (indeed, usually a large number of agents is assumed). Whereas with individual learning it is the same individual who changes his behaviour, with social learning what matters is the evolution of the population behaviour: certain types of beliefs or behaviour might become more common within a population or society. Of course, some people would say that this is not "learning" as such, since learning must involve some mental processes and the mental state of "understanding". However, this is an issue which lies beyond this chapter: I will simply follow the common usage and call all forms of adaptation and selection "learning".

An archetypal example of social learning is Darwinian natural selection. Suppose that we have a particular species: the giraffe. The giraffe develops a long

neck so that it can eat leaves high up on trees⁵. Now, we can view this as a design problem: suppose that you were designing a giraffe. You can make the giraffe have a short neck or a long neck. There is a cost to a longer neck: it reduces mobility, uses up more energy and requires more food to keep it going and so on. It also has benefits: the giraffe has access to leaves that are beyond the reach of other land based animals. The question is whether the costs outweigh the benefits. Individual giraffes never learn about this: they have the neck they are born with and that is it⁶. However, if the marginal benefit of a longer neck outweighs the marginal cost the giraffes with longer necks will prosper and have more children who will tend to inherit the longer neck gene. This will go on until the point at which the marginal cost outweighs the benefits (or so the simple story goes). We can say that although no individual giraffe learns anything (they just hang out, eat leaves and try to avoid being eaten by lions), the giraffe species has “learned” the solution to the problem of neck design. Now, this is perhaps a non-standard use of the term “learn”, but it is one which is standard in this literature, perhaps made more palatable in an economics by the fact that we are not talking giraffes but humans. Darwinian natural selection is an extreme form of social learning: in economic models we might expect individuals to learn within the process of natural selection. There should be an interaction between individual and social learning.

The plan of this chapter is as follows. In section 8.2 we will take a look at Antione Augustine Cournot’s best response framework and the models of evolutionary biology (evolutionary stable strategies and the replicator dynamics). We will discover the close relationship between Cournot’s concept of the Nash equilibrium and evolutionary equilibrium made clear by John Maynard Smith. As I shall argue, the assumption of random matching underlying the biological models are not appropriate for most economic phenomena. In section 8.3 we will examine local interaction models which abandon the random matching assumption and replace it by agents interacting over time in a fixed network of relationships (Donut world). In section 8.4 we introduce an explicitly economic context to the learning process. In the Duopoly Archipelago, there is a whole economy of markets, within each market there are two firms playing some sort of market game. The new feature is that there is a capital market which imposes the discipline on all firms that they earn at least average profits in the long-run: the capital market imposes a selection criterion, survival of those that

manage to keep up with the population average – “keeping up with the Joneses”. There we find the surprising result that the in each market in the economy is driven towards collusion.

8.2 Social Learning: from Antoine Augustin Cournot to John Maynard Smith.

In this section I look at two types learning model. The French economist Antoine Augustin Cournot. (1801-1877) is central to both: in the 1980’s John Maynard Smith extended the empire of economists from rational economic agents to the natural world of dumb beasts. First I will review Cournot’s duopoly model as a learning model. Second, I will review the biological model of evolution with random matching. What we shall see is that there is a very close relationship between evolutionary models and the economist’s concept of a Nash equilibrium. Indeed, in his book *Evolution and the theory of games* the British evolutionary biologist John Maynard Smith showed that we can look at the outcome of evolutionary processes as a Nash equilibrium.

8.2.1 Cournot and Best response dynamics.

Let us start from the beginning, Cournot’s familiar model of duopoly, as we have discussed in the previous chapters on *equilibrium and explanation* and *oligopoly theory made simple*. The process of adjustment to equilibrium involves an alternating move structure: we can think of time divided into discrete periods and firms alternately set their output for the next two periods. Cournot introduced the idea of the *myopic best response* dynamic: the firm that sets its output in period t chooses the best response to the output currently produced by the other firm. Now, we can say two things about this simple “society”. First, we can think of the firms as *learning* about each others’ behaviour: each period they update their own beliefs about the other firm’s output and adjusts its own behaviour appropriately. Also, the process will (under certain conditions) lead the firms to play the Nash⁷ equilibrium in outputs. To see why, recall that any stationary point in this *learning* process occurs only when each firm is choosing a best response to the other firm, precisely as defined in the

notion of a Nash equilibrium. We can see that there is a relationship between learning process and the equilibrium here. The learning mechanism (myopic best response) defines a dynamic process (the time path of outputs). The equilibrium can be thought of as a stationary point in this dynamic process that is stable. The Cournot adjustment process has been the subject of much criticism: why should firms be so myopic? However, in the context of bounded rationality, assuming agents are dumb is not such a bad thing! As we have discussed before, exactly how far to dumb down is a big issue.

8.2.2 Replication is the name of the game: Evolutionary Biology.

Another type of social learning model comes from evolutionary biology. These are not easy to adapt to economic applications: however, since lack of realism has not often deterred economists, let us proceed with the following model, which can be seen as a metaphor or parable. Consider a population of economic agents who each live one period. They are randomly matched with each other each generation. The economic agents have offspring: the payoff of the agent during its lifetime determines the number of its offspring. How will the population evolve over time? Well, we can define an agent by the action it takes (e.g. the level of output it chooses). We can then describe the population at time t by the proportions of each action which prevail at that time: for example, if there are three types of agent $\{A, B, C\}$, then we have the vector of the 3 population shares $[P_A, P_B, P_C]$ with $P_A + P_B + P_C = 1$. From an individual agent's point of view, what matters is the action played by its opponent when it is alive, since this determines its own payoff (it does not care about other agents' payoffs). However, from the point of view of the population, all that matters is how each particular type does: on the assumption of random matching, players of a particular type are evenly spread over the population. For example, if $P_A = 0.3$, $P_B = 0.2$ and $P_C = 0.5$, then if we take type A for example, 30% of type A agents will be playing type As: 20% type Bs and 50% type Cs. This can be represented by the array:

$$\begin{cases} P_{AA} = 0.09 & P_{AB} = 0.06 & P_{AC} = 0.15 \\ P_{BA} = 0.06 & P_{BB} = 0.04 & P_{BC} = 0.1 \\ P_{CA} = 0.15 & P_{CB} = 0.1 & P_{CC} = 0.25 \end{cases}$$

The first row represents the distribution of type A: P_{AA} the proportion of type As matched with type As; P_{AB} the proportion of type As with matched with Bs and so on. The second row represents the distribution of type B agents over the population. With random matching, this distribution is easily calculated: $P_{Ai} = P_A \cdot P_i$ where $i=A,B,C$.

Now, let us suppose that there is a $n \times n$ payoff matrix Π with elements $\pi(i,j)$ which give the payoff to a strategy i when it plays a strategy j . With three strategies we have the 3x3 matrix:

$$\Pi = \begin{bmatrix} p(A, A) & p(A, B) & p(A, C) \\ p(B, A) & p(B, B) & p(B, C) \\ p(A, C) & p(C, B) & p(C, C) \end{bmatrix}$$

For example, the player “type” or strategy might be outputs if the matched players play a Cournot duopoly game. This might not be very realistic: it is hard to imagine firms being randomly matched: one period you play a sock firm, the next a bicycle firm. However, continuing to ignore realism as an issue and for the purpose of exposition, let us suppose that the three types are actually output levels X_i :

Expository Parable of the Randomly Matched Cournot Duopolists.

- When any two individual agents are matched, the Industry demand they face is $P=1-X_i-X_j$.
- There are no costs.
- Type A produces output $X_A=1/2$; type B produces $X_B=1/3$; type C produces $X_C=1/4$.

In this case we have the payoff matrix $\pi(i,j)= X_i(1- X_i- X_j)$ – see *oligopoly theory made simple* for more details- with payoffs both as exact fractions and decimals to 3 places.

$$\Pi = \begin{bmatrix} 1/8 & 5/48 & 1/16 \\ 5/36 & 1/9 & 1/15 \\ 1/8 & 1/10 & 0 \end{bmatrix} = \begin{bmatrix} 0.125 & 0.104 & 0.063 \\ 0.139 & 0.111 & 0.067 \\ 0.125 & 0.100 & 0.000 \end{bmatrix}$$

Clearly, there is a unique strict Nash Equilibrium here: both firms produce $1/3$. This is the Cournot-Nash equilibrium. To see why, let's consider the best response of the row⁸ player i . If firm J produced $1/4$, then i 's best response (look down the first column) is $1/3$ (since $5/36 > 1/8$). If firm j plays $1/3$, the best response is $1/3$; if j plays $1/2$ then $1/3$ is also best. In fact, in this simple example, Strategy B ($1/3$) is a *dominant strategy*: whatever the other player does, an output of $1/3$ yields the best payoff - the second row has the largest element in each column. Of course, we could have constructed things so that there was no dominant strategy, but the types chosen are salient: $1/4$ is the joint-profit maximizing strategy; $1/2$ is both the Stackelberg leaders output and half the Walrasian output. Having outlined the basic structure of the model, we now need to consider the population dynamics: we will take the example of the *replicator dynamics*.

8.2.3 The Replicator dynamics: even educated Flees do it.

The basic idea behind the *replicator dynamics* is simple: strategies that have a higher payoff have more offspring: their share of the population gets bigger. Success breeds success, failures fade away. Let's have a quick look at the mathematics. For those of you who do not like equations, just skip the rest of this section and move straight to 2.4.

The average payoff of strategy i at time t ($\Pi_i(t)$) is defined as the weighted sum of its payoffs playing each strategy, where the weights are the population proportions. Hence for strategy A

$$\Pi_A(t) = P_A(t)p(A,A) + P_B(t)p(A,B) + P_C(t)p(A,C)$$

whilst the *average payoff* over all firms at time t is the weighted average of the payoffs off each strategy over the whole population

$$\Pi(t) = \sum P_i \Pi_i(t) = P_A(t)\Pi_A(t) + P_B(t)\Pi_B(t) + P_C(t)\Pi_C(t)$$

We can now model the process of evolution. The simplest form of evolution is to suppose that the population dynamics are given by the *replicator dynamics*:

$$g_i(t) = \frac{P_i(t+1)}{P_i(t)} = \left[\frac{\Pi_i(t)}{\Pi(t)} \right]$$

The growth of the proportion of type i is equal to the ratio of its payoff to the population average. The point here is that the proportion of a particular type increases (decreases) in proportion to the extent that its payoff is above (below) average. There is a simple story underlying this: the number of offspring is a linear function of the actual payoff⁹.

8.2.4 Alien Invasions and the Evolutionary Stable Strategy (ESS)¹⁰.

John Maynard Smith, the British evolutionary biologist, introduced the concept of the evolutionary stable strategy (ESS). An ESS strategy is one which is stable if there is a small invasion by another strategy. Suppose the whole population is playing one strategy. Now let a small ε -invasion happen: an ε -invasion occurs when a proportion of size ε invades the population. The initial strategy is ESS if the ε -invasion will not succeed - it will die out. This can be expressed formally in the following way: suppose that start from a situation where all of the population (earth people) is playing some strategy i : from our example, i can be one of $\{A, B, C\}$. The payoff of all firms will then be $\pi(i, i)$. Now, if a proportion of players of type j (other than i) invade the population, the average payoff of the alien invaders will be $(1-\varepsilon)\pi(j, i) + \varepsilon.\pi(j, j)$. The alien invaders are almost certain (with probability $(1-\varepsilon)$) to meet someone playing strategy i ; with a small probability ε they meet one of their own¹¹. Likewise, the earth people playing i will have the payoff $(1-\varepsilon)\pi(i, i) + \varepsilon.\pi(i, j)$. The condition for strategy i to be an ESS is in maths:

Definition 1: Strategy i is ESS if for all j other than i

$$(1-\varepsilon)\pi(i, i) + \varepsilon.\pi(i, j) > (1-\varepsilon)\pi(j, i) + \varepsilon.\pi(j, j).$$

In plain English, the alien invaders earn (strictly) less than the earth people. The left-hand of the inequality is the payoff of the earth people; the right-hand the aliens. Hence if population growth depends (positively) on payoff, the aliens will die out. Now, we come to an amazing result. First we have to understand the notion of a *strict* Nash equilibrium. A strict Nash equilibrium occurs when the equilibrium strategy

yields *strictly* more than any other possible strategy: a *weak* Nash equilibrium occurs when the equilibrium strategy earns no less (i.e. *weakly* more) than any other possible strategy. In particular, a *sufficient* condition for strategy i to be *ESS* is that it is a strict Nash equilibrium strategy. In fact John Maynard Smith (1982) showed that the above definition of an *ESS* was equivalent to the following:

Definition 2: strategy i is *ESS* if

- (a) $\pi(i,i) \geq \pi(j,i)$ for all j other than i .
- (b) if $\pi(i,i) = \pi(j,i)$, then $\pi(i,j) > \pi(j,j)$.

As we can see, part (a) of the definition is simply the standard definition of a Nash equilibrium. If the Nash equilibrium is strict, then it is automatically an *ESS*. If we have a non-strict Nash equilibrium, we need to have the additional condition (b): the alien invaders do worse against themselves than the earth people. This result has the amazing implication that we can use game theory to model evolutionary biology! I recall in 1982 dining at Christchurch college Oxford. Neither the fact that I was sat next to an elderly cleric called a “cannon”, nor the fact that the food was cold by the time it had reached the high-table from the kitchen were the most amazing thing that evening. No, I was most surprised by a zoologist who told me that he was applying game theory to animal behaviour. What seems a commonplace now seemed amazing then, since we all used to look at game theory in terms of rationality wonderland. Well, as the evening wore on (and after more glasses of wine and surreal conversations with the elderly cannon) it seemed pretty sensible.

However, the final relationship is between the replicator dynamics and the *ESS* concept. Again, there is a strong relationship between the two ideas: *every ESS is an asymptotically stable steady state of the replicator dynamics*¹². A steady state is a state which is unchanging over time: in this case we can think of the state being the vector of population proportions. A steady state is *asymptotically stable* when the system returns to the steady state whenever there is a small deviation from equilibrium¹³.

What is the relationship between a Nash equilibrium and the replicator dynamics? Well, any steady state that is asymptotically stable under the replicator dynamics has to be a Nash equilibrium. This is both important and obvious. A Nash

equilibrium strategy has to be a best response to itself: this is also a necessary condition for the replicator dynamics to be stable around a steady state. To see why, suppose that a strategy was not the best response to itself: in terms of our example, let us suppose that we have a steady state where all firms are “collusive” type A’s. In this case, suppose that we move away from this a little and introduce some Cournot type B’s. From the payoff matrix the type B’s will earn more than the type As, and so the proportion of type B’s will increase, leading to a move further away from the initial steady state. This argument will hold for any non-Nash equilibrium Strategy.

However, whilst all stable steady states are Nash equilibria, not all Nash equilibria are stable. For example, let’s augment the strategy space to include a type *D*, which always produces 1 unit of output. This strategy yields a zero payoff for itself and any strategy it plays against¹⁴. In effect, the price is kept at zero whatever the opponent does: it seems appropriate to name it the “Bertrand” or *party pooper* strategy. This strategy is a (weak) Nash equilibrium, the Bertrand equilibrium. However, it is certainly not stable: suppose that some collusive firms invade. These may earn zero most of the time when they play Bertrand firms: however, when they meet each other they earn a positive profit, so that they will thrive and increase in number, whilst the Bertrand firms decline.

The relationship between the three concepts of Nash equilibrium, stability of the replicator dynamics and ESS for steady states is depicted in Figure 1. We have concentric circles: the largest set is the set of all Nash equilibria; within that we have the set of stable steady states; within that is the set of *ESS*; within that is the set of strict Nash equilibria¹⁵.

Fig 8.1

We have come full circle: we started with Cournot and his equilibrium. He saw the equilibrium outcome as resulting from a dynamic process of adjustment: the steady state arising out of it. The resultant Nash equilibrium has formed the basis for imperfectly competitive models, from Edgworth’s price-setting duopoly model (1889) and the Robinson/Chamberlin model of monopolistic competition (1933) to the present day. We have also seen the same equilibrium concept playing a crucial role in evolutionary biology. In between, we have the more orthodox perspective of super-rational agents with common knowledge playing games. It is amazing that the same equilibrium concept can be seen as arising from such different processes and perspectives.

8.2.5 Random Matching and Economics.

The evolutionary models used in biology are perhaps not well suited to economics. Most of these are based on the crucial assumption of random matching. Most economic interactions are repeated. We buy and sell with familiar traders over time. We work for the same firm, buy from the same shops, visit the same restaurants. This is the same whether we think of the household or the firm. If we are thinking of oligopoly or collective bargaining, then random matching is particularly inappropriate! In economics, the modelling of evolutionary forces by such biological models might be thought to be inappropriate except for special cases.

However, whilst I would myself counsel strongly against the unthinking and literal use of such biological models, we can think of the biological process not in terms of its microfoundations, which are inappropriate, but rather as a *metaphor*. The *evolutionary metaphor* merely says that forms of behaviour (strategies) that are more successful (earn higher payoffs) tend to become more common. That having been said, there can be no substitute for an appropriate framework for modelling economic and social interaction.

8.3 Social learning in Human Societies: Gabriel Tarde

*Whatever a great man does, the
very same is also done by other men.
Whatever the standard he sets,
The world follows it.*

Bhagavad Gita, 3.21.

There are powerful forces of learning in human societies that are not captured in the basic natural selection model. This was recognised by the French social theorist Gabriel Tarde (1843-1904). He was a lawyer and judge who for obvious reasons thought a lot about the causes of crime. He developed some general principles which he called the *laws of imitation*. He thought that people learn from one another through a process of imitation, and that activity or behaviour seen in others tends to reinforce

or discourage previous habits. He also observed that the process of diffusion in human society often follows an “S-curve”, otherwise known as the *logistic curve*, as depicted in Figure 2.

Figure 8.2

What happens in the S-curve is roughly as follows. Someone has an idea: let us take the concrete example of a new method of breaking into a house. At first, only that person knows about it, plus possibly a few close friends whom he tells about it (possibly when he/she is in prison). But these friends can tell their friends and so on: the process of growth here is exponential: each new person who catches on to the idea can pass it on to a few others. This explains the initial convex part of the curve: the absolute number of new people adopting the innovation in each period (the slope of the curve) is increasing up to time T' . This process cannot go on forever, however, since the population is finite! What happens eventually is that a saturation point is reached. Eventually when a new person learns of the idea, they will find that most people they tell the idea to will already know it. The process of growth will thus slow down and possibly there will remain some people (e.g. non-criminals) who will never adopt the new technique for housebreaking. After time T' the number of people adopting the innovation slows down and the curve becomes concave. Of course, ideas come and go and the world does not remain still. If lots of criminals adopt the new technique of housebreaking, then the police and security firms will develop counter measures which house owners will start to adopt. After a period of time the new technique will become less useful and there may be a period of decline. In ancient Greece, one method of housebreaking was to tunnel through the walls of the house. This method relied on a mud-wall construction, and the technology died out when this construction technology became less common.

This theory of social diffusion has been widely developed and applied in a variety of contexts. In particular, it provides one of the basic models for marketing: firms are keen to look at ways of speeding up the process of adoption of a new product and extend the life of an existing one (see for example Kotler 1986). It is also used by economists as a model for the diffusion of technical progress (the path breaking paper here was Mansfield 1961), applied in health (the theory of diffusion of medical practices and diseases, Coleman 1966) amongst others.

Certain factors have been identified as important in the spread of an idea. For example, the adoption of the idea by opinion leaders can be crucial: if a widely known and respected individual is known to adopt an idea, it gives others the inspiration and confidence to try it out. A firm may not risk trying out a new technology until it has seen that some of the large established players have taken it seriously. We can think people having an agenda: these are the ideas or actions that people take seriously in the sense that they might actually think about adopting them. Because of the limitations of bounded rationality, people do not think about everything all of the time: they only think about a few things most of the time. We all know this from our own experience. We know that certain types of food and drink are bad for us: however, although we know and are aware of the healthier alternatives, we still end up eating the same old food most of the time. It takes some effort to change habits, to put new ideas (in this case a new diet) onto our agenda, so that we think about them seriously when we take decisions. Seeing someone whom you respect or identify with in some way adopting the idea is a way of putting it on your agenda, which makes it more likely that you will adopt it. This was exactly Krishna's argument to Arjuna quoted in the *Bhagavad Gita*.

8.3.1 Welcome to Donut-world.

One way of thinking about the process of interaction is to imagine society as a *donut*. A donut is a three-dimensional Torus: a Torus is a network without any edges. To make this clear, think about network consisting of houses and paths. In each house there lives an economic agent. The houses are connected by paths¹⁶. We can represent a society by a map of the houses and paths, as in Figure 8.3. Now, houses and paths can in theory be built anywhere: however, we can imagine that planning laws dictate a particular structure called a *lattice* or *grid* as in Fig.8.3. In a lattice, the houses are built in equally spaced rows and columns, whilst the house is connected to other houses by paths which are either East-West or North-South. Thus a house is only connected to its 4 immediate neighbours (going clockwise and starting at the top, North, East, South and West): it is not connected with its other 4 neighbours (who are Northwest, Northeast, Southeast and Southwest). We can think of the *neighbourhood* of the agent: these are the other agents with whom the agent interacts directly. The neighbourhood is defined by a number r : this is the number of paths the agent can

travel to interact: if $r=1$, then the neighbourhood of the agent consists of its 4 immediate neighbours. If $r=2$, then the neighbourhood expands to include 8 other houses (a total of 12). Lets keep life simple, and suppose that $r=1$: the neighbourhood consists only of the folks next door.

Figure 8.3

In Figure 8.3, there is an edge to the lattice, where the houses stop. If you live in the middle of the page, you will have the regular 4 neighbours. However, if you live on the edges, you will only have 3 neighbours: if you live at the corners, you will only have 2 neighbours. Now, lets talk donuts. First forget the origami: *do not attempt to tear out the page and fold it at home* (it won't work¹⁷). In your mind, consider what would happen if you joined up the top and the bottom row together and also the left and the right. Think about it for a while: the end result would be a donut, a surface without edges or corners. This is a very useful concept, since it means that every house is the same: all houses have the same number of neighbours. A one dimensional Torus can be represented in two dimensions as a circle¹⁸: the two dimensional Torus can be represented in three dimensions as a donut. The important thing is that the dimension of the Torus is one less than the dimension it occurs in (much the same as the surface of the earth is two dimensional¹⁹ but occurs in three dimensional space). Rather than trying to draw a real donut, we can represent it by imagining that in fact there are paths going from all of the houses on the left side of Fig 8.3 to the houses on the right side, and those on the top to those on the bottom (corner houses would thus have two new paths). Maybe one day fast food outlets will sell "flat pack" donuts which you assemble before eating.

So, here we have our simple society. Let us suppose that each household is growing food (indeed some of the earliest studies of diffusion were in agriculture, Ryan and Gross 1943). Each agent can see the gardens of the houses in his neighbourhood and the methods of gardening; Now, suppose that one household innovates in period 1: it works out that if it rotates the crops in a certain way disease is reduced and output increases. We can take the idealised case first: suppose that neighbours see exactly what is going on and will always adopt a new technology with certainty if it is beneficial. In this case, in period 1 all neighbours will have seen what was done in period 1 and also that it yielded a greater harvest. So, in period 2 they will do the same thing: there are now 5 households rotating the crops. In period 3 their

neighbours will also do the same thing: and additional 8 households bringing the total to 13. Now, suppose that this is an infinite lattice (it goes on forever): then the growth will result from each new house in period t generating $4(t-1)$ new houses in the next period. The sequence with one house starting is thus: 1, 5, 13, 21, 37..... The total number of houses with the new technology at time t , denoted $H(t)$ is thus given by the recursive relationship $H(t)=H(t-1) + 4.(t-1)$, along with the assumed initial value $H(0)=0, H(1)=1$.

Now, let us assume that we are in Donut world, in a 5×5 3-D Torus. In this case we have the constraint $H(t) \leq 25$. What will happen? Well, for the first three periods, everything is as in the infinite lattice case: $H(1)=1, H(2)=5, H(3)=13$. Now, in period $t=4$, only 8 new houses adopt the innovation: the 4 houses at the “edge” of the square in period $t=3$ are next to houses that have already adopted the technology. Hence $H(4)=21$. The 4 households at the “corners” of the square are not reached until the next period: in $t=5$ there are 4 new houses adopting the new technology, so that $H(5)=25$, and $H(t)=25$ for $t \geq 5$. In Figure 8.4 we can plot the diffusion of the technology: it indeed follows a roughly S-shaped curve: the increases in absolute terms are: 1,4,8,8,4,4.

Fig. 8.4

This simple story is deterministic. Lets introduce some uncertainty or randomness into it. For example, suppose that households will not be looking at each others gardens all the time and may (due to fog or rain) not observe them with great accuracy; the output of the farms has a random element; the households that see a high output might not bother to adopt the new technology (due to inertia or laziness). Let us suppose that we start from a situation where all gardeners are doing the same thing: they will obtain the same yield as each other, subject only to a random element (the luck of the draw) each harvest. This randomness means that the actual path of diffusion will be random, depending on what happens. For example, the diffusion might take some time to get started: the household with the new method might be unlucky for a few periods and its output might not be particularly high; even if it is high, the neighbours might not notice; even if they do notice they may not do anything about it straight away. *However, the important thing to note about randomness and uncertainty is that if there is a enough time, then everything that can*

happen will happen. From the perspective of eternity, everything is possible. The exact timing of events is random: but as the time available becomes longer and longer, even quite unlikely things can happen. It is like throwing dice: it is fairly unlikely that you will throw a double 6 in any one throw, but as you keep throwing the event becomes more and more likely. Throw the dice a thousand times and it is almost certain you will throw a double six at least once. This is why, in social learning models researchers often concentrate on the asymptotic or long-run properties of the system (modelled mathematically as what happens when as t tends to infinity). In terms of social learning, the effect of the randomness is merely to slow up the path of diffusion and make the exact path and timing uncertain. The end result will not necessarily be changed (see for example Bala and Goyal 1998 for an analysis of learning from neighbours with local interaction with an explicit model of learning).

8.3.2 Diffusion in a strategic context.

The case of an innovation is *non-strategic*: my method of cultivating vegetables does not affect yours²⁰. Now, let us think of a strategic interaction, where I actually undertake some sort of economic activity with my neighbours. For example, consider the Prisoner's Dilemma PD . There are two strategies: cooperate C or defect D . The payoff to the farmer of playing strategy i against j $\pi(i,j)$ are as follows: $\pi(C,C)=2$, $\pi(C,D)=0$, $\pi(D,D)=0$ and $\pi(D,C)=a$ which we represent the payoff matrix

$$\Pi_{PD} = \begin{bmatrix} 2 & 0 \\ a & 1 \end{bmatrix}$$

Each farmer plays the PD all of his 4 neighbours. However, he²¹ cannot customise: at time t he can only play one strategy with each neighbour, either C or D : one size must fit all comers.

Now, clearly, since D is a dominant strategy, the issue might seem pretty trivial: choose D , since it is the best strategy whatever the competition does. However, if everyone had that attitude all game theorists would be unemployed. So lets assume that things happen differently. Each farmer does what he does. However,

he observes the payoffs of his neighbours. If a neighbour is doing better than him, he will imitate the neighbour: if more than one is doing better, he will imitate the one with the highest payoff. Lets call this process “*imitate your best neighbour*”²². Now, the payoff of a particular agent will be the sum of the payoffs he earns form his 4 neighbours. If we take the case of the *PD*, we then have a variety of possibilities: for a given strategy chosen by the farmer Giles, there are 4 possible combinations of strategies he can face: {C,C,C,C}, {C,C,C,D}, {C,C,D,D}, {C,D,D,D}, {D,D,D,D}. Since there are two possible strategies farmer Giles can choose, we can represent the payoffs of the farmer in table, with $a=2.5$

	C	D
CCCC	8	10
CCCD	6	8.5
CCDD	4	7
CDDD	2	5.5
DDDD	0	4

Table: Payoffs and the Neighbourhood strategies with *PD* ($a=2.5$).

Now, consider what might happen here. Let us do a few thought experiment in Donutworld. What will the learning rule “*imitate your best neighbour*” generate at the social level?

Case 1: Suppose the Torus is an even-numbered square (for concreteness a chess board, 8x8). 50% of farmers choose *C*, and 50% choose *D*. Furthermore, suppose that like a chess board, the *C*s and *D*s alternate. Every farmer will have two *C*s and 2 *D*s in his neighbourhood. Looking down the table, the *C* farmers will be earning 4; the *D*s 7. In this case, the *C* farmers will look enviously at their *D* counterparts, and imitate them. So, the next period all firms will choose *D*, which is of course the Nash equilibrium of the *PD*. The equilibrium under the “*imitate your best neighbour*” is a steady state where all farmers adopt *D*.

Case 2: As in case 1, but the *Cs* and *Ds* are partitioned into two separate blocks. The top half of the donut is all *D*, the bottom is all *C*. There are 4 different payoffs here. The *D* surrounded by *Ds* earns 4; the *C* surrounded by *Cs* earns 8. The more interesting case are the 2 borders, where *Cs* meet *Ds*. Given that the Torus is an even square (8×8) these are a straight lines. In this case, each borderline *C* will have 3 fellow *Cs* and one *D*, hence earning 6; each borderline *D* will have 3 *Ds* and a *C*, earning a total of 5.5. Thus, in the next period, all of the borderline *Ds* will switch to *Cs*. Each period, cooperation will spread two more rows²³, until the whole of Donut world is playing *C* after two periods. Again, we have a steady state equilibrium, but with all farmers adopting *C*, the opposite outcome to case 1.

Fig.8.5

Case 3: As in cases 2, except that the square Torus is of odd size (e.g. 7×7) depicted in Fig 8.5. An exact 50/50 split is not possible here. In this case, the top three rows all *D*; the bottom three rows all *C*; the middle row will be a mixture of *Cs* and *Ds*. Now, let us take the case where there are 24 *Ds* and 23 *Cs*: the middle row will have 4 *Ds* and 3 *Cs* alternating. In effect, the middle border between the *Cs* and *Ds* is a zig-zag. Every border *C* in row 4 will have 3 *Ds* and one *C* as neighbours, thus earning 2; 2 of the border *Ds* have 3 *Cs* and a *D* as neighbours and earn 8.5; the end two have two of each and earn 7. Clearly, all of the *Cs* in the mixed border row 4 will switch to *D*. There are also 4 *Cs* in row 5 who have one *D* neighbour in the mixed row: they will be earning 6, and hence will also switch to *D*. Hence in period 2, row 4 has become all *D* and row 3 becomes mixed (the same as row 4 in the previous period). At the same time, there is a straight border between row 1 (all *D*) and row 7 (all *C*). As in case 2, the *Ds* in row 1 will all switch to *C*. Hence, the state of the Donut economy in period 2 is the same as period 1, except that it has been “rotated”: the borders move “up” one each period (like an Escher figure, going up from row 7 means going to row 1). Thus any particular farmer will spend 3 (or 4) periods *D* and 4 (or 3) *C* depending in which column he finds himself. This is the attractor of the imitation dynamics: it is not a steady state, but a cycle: it is rather like the human wave of hands that passes around the football stadium as people raise their arms up if the person next to them does.

What these three thought experiments show us is that in the world of local interaction, there is no inevitability about the Nash equilibrium coming about, even in the stark case of the PD where there is a dominant strategy. We can get cycles (case 3), or convergence to uniform populations of either all C (case 2) or all D (case 1). The history depends very much on the initial conditions and the exact structure of the payoff matrix. For example, the larger $p(D, C)$ (i.e. a), the less likely is C to survive; also, the C s need to live together and apart from the D s to survive.

In this section we have seen how social learning can be modelled in both a strategic and a non-strategic setting. Essentially, we can represent social interaction as a network (in the case of Donut world, a lattice Torus): agents repeatedly interact within a neighbourhood. Clearly, Donut world looks a bit more like a real economy than the world of random matching. However, lets go a step further and try to construct something that looks even more like an economy: time to move on to the *Duopoly Archipelago*, a place where all aspirations are met in the long run and everything is possible for he who decides to experiment.

8.4 Economic Natural Selection: Keeping up with the Joneses.

"The best monopoly profit is a quiet life" John Hicks (1935).

"This is the criterion by which the economic system selects survivors: those who realize positive profits are the survivors; those who suffer losses disappear" Armen Alchian (1950, p.213).

The idea of Natural selection in economics is not new. It has long been argued that firms must earn at least normal-profits to survive in the long-run²⁴. Failure to achieve this will activate some market mechanism which will lead to the ownership or the control of the firm changing. These mechanisms include:

- Bankruptcy: the firm becomes insolvent and is forced to stop trading. Its assets are then sold off.

- The shareholders replace the existing of managers.
- The shares of the firm are purchased by the managers of another firm who replace the existing managers.
- Debtors are able to reschedule outstanding debts and impose changes on the firm.

In general, we can think of the mechanisms as reflecting the operation of the capital market in its widest sense. The performance of a particular firm is measured against the performance of other firms. The ultimate bottom line for the capital market is the profitability of the firm and its ability to deliver dividends to shareholders and/or keep up scheduled loan repayments. An extreme form of failure is insolvency or bankruptcy which occurs when a firm is unable to cover its expenditure with current income: the cash coming is less than the liabilities it is incurring. It is against the law to continue to trade when insolvent: when a company is insolvent an outside agent is called in to take over responsibility for the company (in Britain, this person is called *the receiver*). The decision may be taken to liquidate the company: i.e. sell off its assets and meet as many of the outstanding liabilities as possible. Alternatively, the decision may be taken to find new managers to continue running the company as a going concern.

However, even if managers of the firm are making a profit and have no cash flow problems, there are still constraints. There are a variety of benchmarks against which they are judged by the capital market. In the first instance, the benchmark is provided by similar firms in the same or related lines of business. If firm X and firm Y are in the same industry, their profitability (rate of return on capital) should be the same: if firm X consistently under-performs relative to Y, then this is a good indicator that the strategy of X is not the best. However, in the long run, there is an *arbitrage condition*: the rate of return must be the same for all possible investments²⁵. The argument here is the same as all arbitrage arguments: if your capital is earning less in one place than another, then shift it to the place earning more. So, the capital market links together all of the firms in the Duopoly Archipelago. Be they selling Pizzas, making air conditioners, or an airline, the capital market evaluates them and reduces them to the same thing: money making machines. *The capital market requires them to be equally efficient money making machines.*

The capital market reflects the aggregate performance of the economy as represented by average profitability. In this paper the level of normal-profits is taken to be the average level of profits in the economy and explores the implications of this hypothesis in the context of an economy consisting of many oligopolistic markets. Under fairly general assumptions there are powerful long-run forces pushing the firms in each market towards collusion. What differentiates the approach here is that the evolution of the economy is inherently social, in that it is the level of average profits in the whole economy over time which drives the behaviour of firms.

In this section I model the behaviour of firms using an aspiration based model of bounded rationality. *The key feature of this model is to link together the aspirations of firms with the level of normal profit by requiring that in the long run the aspiration level of all firms is to have at least normal profits.*

8.4.1 Welcome to the Duopoly Archipelago.

Imagine an archipelago of islands: each island represents a market. On each island there are two firms. The markets and firms on each island are the same in terms of size, costs and so on. We can picture the economy in terms of each island being having two houses (firms) linked by a single path. The neighbourhoods of each firm consist only of its competitor on that island (the economy is not directly interconnected as in Donut world).

Firms have a finite strategy set with K pure strategies $i, j = 1 \dots K$. For concreteness, we can think of the strategies as output levels X_i as previously, with no cost and linear demand. We need assume very little about the structure of the payoff matrix Π of the constituent duopoly game, except that the joint-payoff can be maximized by a payoff-symmetric outcome. An outcome can be thought of as a pair of strategies (i, j) : it is payoff-symmetric if the payoffs are the same for both firms: $\pi(i, j) = \pi(j, i)$. Clearly, the leading diagonal of the matrix Π is payoff symmetric: however, it is possible in general that off diagonal terms might also be payoff symmetric. In the case of outputs as strategies, the payoff symmetric outcomes will consist exclusively of the leading diagonal. We will therefore assume for simplicity

that equal profits for the duopolists on a particular island means that they are producing the same output ($\pi(i,j) = \pi(j,i)$ implies $i=j$).

We will make the following assumption about the payoff matrix. The state of a market is fully described by the pair of strategies chosen by the firms in that market: which firm chooses which does not matter (except, of course for the firms concerned!). Suppose we are free to choose any pair of outputs $\{i,j\}$: the joint profit maximizing pair(s) S is (are) the pair(s) that maximize the joint profits of the firms, with the maximum joint profits denoted JPM :

$$JPM = \max_{\{i,j\}} \frac{p(i,j) + p(j,i)}{2}$$

In the case of the simple Cournot oligopoly, S consists of the unique pair of outputs $(\frac{1}{2}, \frac{1}{2})$, each firm producing half of the monopoly output: the JPM profits are then $\frac{1}{4}$ (each firm earns $\frac{1}{8}$). Now, let us assume that we maximise joint profits, restricting ourselves to cases where both firms produce the same output:

$$SJPM = \max_{i=1 \dots K} \frac{p(i,i)}{2}$$

Now, the assumption we need to make is that the joint profits are at their maximum when payoffs are symmetric. One way of saying this is:

Assumption 1: $JPM = SJPM$.

Clearly, this assumption is satisfied by the simple Cournot model we are using as an example. If we consider the *Prisoners Dilemma* (PD), we have

$$\Pi_{PD} = \begin{bmatrix} 2 & 0 \\ a & 1 \end{bmatrix}$$

There are two strategies: cooperate C or defect D. $\pi(C,C)=2$, $\pi(C,D)=0$, $\pi(D,D)=0$ and $\pi(D,C)=a$. For this to be a PD we require $a>2$: it must pay to defect when the other person is cooperating: it also ensures that D is the dominant strategy. However, assumption 1 will only be satisfied if $a\leq 4$. To see why, note that if $a>4$, then

$$JPM = \frac{\mathbf{p}(C,D) + \mathbf{p}(D,C)}{2} = \frac{a}{2} > SJPM = 2$$

Hence, for the model to apply to the PD we need to assume that $a\leq 4$.

When we look at the economy as a whole, we can summarise what it looks like in terms of the competition in each market. One way to do this is to take each pair $\{i,j\}$ and measure the proportion of markets (islands) which have firms playing this pair of strategies²⁶, $P(\{i,j\})$. Let $P(S)$ be the proportion of markets where the firms are producing collusive outputs and hence earning the JPM profits.

So, here we have the Duopoly Archipelago. On each island we have a duopoly of firms choosing an strategy pair. We can describe the economy at any time t in terms of the proportions of markets having each possible pair. As a last point, we have to think of the average profits in the whole economy. This is simple to compute: we merely take the combined profits earned with each strategy pair $\{i,j\}$ and then take a weighted average with the population proportions $P(\{i,j\})$ as the weights. The average profits in the economy at time t are then

$$\bar{\Pi}_t = \sum_{\{i,j\}} P_t(\{i,j\}) \left(\frac{\mathbf{p}(i,j) + \mathbf{p}(j,i)}{2} \right)$$

8.4.2 Aspirations in the Duopoly Archipelago.

The concept of an aspiration level has been around for a long time. It has been put forward both as a good model of individual decision making in the mathematical

psychology literature ((Lewin (1936), Siegel (1957)) and as a model of *organisational* decision making with relevance to the firm (Cyert and March (1963), Kornai (1971) and Simon (1947)). Although there are variations, the core idea is simple enough. When attempting to solve a problem, agents (let us think of these as firms) formulate a target: if they achieve this target they will probably stop searching. The aspiration level is a target to which the managers aspire and towards which they plan. As such, the aspiration level is a search heuristic a bit like a stopping mechanism, as for example the reservation wage. In the optimal search literature, the unemployed worker (for example) follows the rule: search until you receive an offer greater or equal to the reservation wage. In fact, under various assumptions, one can derive this as an optimal stopping rule. The aspiration level is a target outcome: if the target is attained by a particular solution or action, then this plan is deemed acceptable and the search is stopped. Of course, aspiration levels can be adjusted in response to experience of the decision makers themselves and outside events. The literature on aspirations does not conceive of the aspiration levels coming from some optimising process: rather it is a boundedly rational attempt to find a good solution.

The aspiration level here has two elements. First, there is the aspiration level as representing *external* forces imposed upon the firm or managers from outside (i.e. the capital market). This is represented by the role of average profitability as an external benchmark for “normal” profits. Second there is the *subjective* element inside the corporate mind, the targets that come up from the interpersonal interaction of managers and others within the firm. In this model, firms at any time adopt a pure-strategy²⁷. Each firm follows the following simple learning rule. It has an aspiration level $\alpha(t)$. If it is earning less than $\alpha(t)$, then it decides to *experiment* with probability 1; if the firm is earning at least $\alpha(t)$, then it will continue with the existing strategy (this is Hick's "*quiet life*" alluded to in the above quote – if it aint bust, then don't fix it).

For simplicity, we assume that all firms share the same aspiration level²⁸, with the aspiration level satisfying the condition that in the long-run it has to be no less than average profits. This seems a reasonable assumption reflecting the role of capital markets in industrialized economies. This means that the aspiration level is

endogenous (as in Borghers and Sarin (1997), Karandikar et al (1998), Palomino and Vega-Redonodo (1999)), reflecting the past and the current profitability of the economy. his might well reflect form specific factors and the history of the individual firm. The assumption we make is that whatever other internal or external factors there are, in the long run the capital market must be satisfied. In fact, we do not have to make explicit the mechanism generating aspirations: we merely impose the following conditions on their evolution (sorry if it looks too technical: just jump to the next paragraph if you do not like the look of it).

Assumption : *Aspirations*.

$$(a) \lim_{t \rightarrow \infty} [\mathbf{a}_t - \bar{\Pi}_t] \geq 0$$

$$(b) \mathbf{a}_t \leq JPM$$

Part (a) says that in the long-run (as t tends to infinity), the aspiration level α_t must be at least equal to average profits $\Pi(t)$. Part (b) also says that firms must not be overoptimistic: the highest realistic aspiration for the firms is that they can earn the *JPM* profit.

Well, we can make things *really* simple: this assumption is satisfied if in each period the aspiration level equals the current average profitability

$$\mathbf{a}(t) = \Pi(t)$$

Well, the aspiration level model here gives the mechanism determining the experimentation by firms. *Experimentation* here means that the firm tries to alter strategy. The question we next ask is: what determines the probability that the firm switches from its existing strategy to another, the *switching probabilities*.

The actual switching probability may be determined by many things: the experience of the firm, what other firms are doing (through *imitation*), or by strategic considerations (as in *best response* dynamics). There might also be some randomness or “noise” in the switching process: mistakes are made, or policies improperly

implemented and so on. This might be very complex to model in detail. However, we make the following general assumption about switching probabilities:

Assumption: *switching probabilities.* There exists some $\gamma > 0$ such that all switching probabilities exceed γ .

What this means is that *anything is possible*: there is a small but strictly positive probability γ that the firm will choose any particular strategy. Of course, some strategies might be much more likely to be chosen: however, no strategy is ruled out. This is not as odd as it sounds: firms sometimes do things that seem pretty stupid with hindsight but looked good at the time!

Whilst we have interpreted switching behaviour as the same firm in two periods changing behaviour, the formal model would be exactly the same if we think of a different firm in each period. For example, a firm in a particular market might exit (due to bankruptcy or death). In this case the switching probability would pertain to the "place" of the firm: the probability that next period the firm taking the place of the existing firm would play a particular strategy.

8.4.3 The Evolution of Collusion in the Archipelago Duopoly.

So, we have set up this archipelago economy, describing the nature of markets (summarised by the payoff matrix) and the behaviour of firms (aspirations and switching probabilities). What happens to it? Well, this is easy to describe. Let us take for simplicity the simplest case where the aspiration level at any time t equals the average profitability, $\mathbf{a}(t) = \bar{\Pi}(t)$. Hence, at period t , we can divide market/islands into three categories:

- *Above aspiration:* Both firms are earning at or above the average. If both firms are above aspiration, then they will just keep on doing what they are doing.
- *Below aspiration:* both firms are earning below average profit. both will experiment: under the assumption that switching probabilities are all strictly

positive, if both firms are experimenting, then anything can happen! There is a strictly positive probability that any pair of outputs/strategies will be chosen.

- *Mixed: one firm above, one below.* In this case, the firm that is meeting its aspiration keeps on with its existing strategy: the one that isn't experiments.

Now, the exact evolution of this economy will be quite complex and depend on the exact switching rules used etc. However, *we can say something in general that will hold for all archipelago economies that satisfy the 3 assumptions we have made:* namely (a) that the payoff matrix has the property that joint profit is maximised with equal profits for both firms; (b) that aspirations tend to average profit in the long run; (c) that when experimenting, anything is possible.

First, consider any industry where both firms are choosing the collusive strategy and earning JPM . Clearly, it is never possible for average profits to exceed JPM : $JPM \geq \Pi(t)$. Whilst it is possible for an individual firm to earn in excess of JPM , it is not possible for two firms in the same industry, nor for all firm in the economy. Hence, industries that are collusive will necessarily be in the above aspiration category. Furthermore, once an industry arrives at collusion it will stay there forever! In technical terms, this is called an *absorbing state*: once you arrive in this state, you are "absorbed" and never leave it. It is a bit like the cockroach motel: the roaches check in, but never check out. An astronomical analogy is a black hole: matter goes in but never comes out again²⁹. *So, over time, we can be sure that the proportion of industries in the economy which are collusive will never get smaller: it must either grow or at least stay constant.* In the case of Cournot duopoly, the collusive outcome involves both firms producing an output of 0.25 and earning 0.125 each.

Second, let us consider the case of industries where both firms are below aspiration. These will tend to be competitive industries, where both firms are producing a large output and earning low profits. Both firms will be experimenting: hence literally any outcome (i.e. strategy pair) is possible, including the collusive outcome. There is a strictly positive probability that both firms will choose the collusive outcome. Looking at the economy as a whole, we will observe a strictly

positive flow from those industries that are below aspiration into the collusive absorbing state.

Let us put these two facts together: once you become collusive, then you remain collusive; a proportion of below aspiration industries become collusive. In the end, if we look at the long run of the economy, we can see that the proportion of firms below aspiration must eventually disappear: they will be absorbed by the collusive state that is in the above aspiration category. This means that in the end, there can only be two categories of firms left: the above aspiration and the mixed.

Third, consider the mixed aspiration category. The mixed aspiration category cannot survive in the long run either. To see why, just note that there is a continual flow of industries from the mixed to either the above aspiration or the below aspiration categories. To see why, note that under the assumption that all switching probabilities are strictly positive, there is a positive probability that the below aspiration firm in a mixed industry will choose the same strategy as its competitor. If this happens, then the next period both firms will be earning the same profit and hence be either both above or both below aspiration. Since the proportion of industries with both firms below aspiration must go to zero in the long run, this flow from mixed industries must in the long run be to above aspiration industries, with the proportion of mixed industries going to zero.

Finally, consider the above aspiration category excluding the collusive industries. Since the proportions of firms with one or both firms below aspiration goes to zero over time, it follows that all industries must be above aspiration. However, how can all firms be at or above the average? Well, there are two ways. First, all industries arrive at the situation where they are all in the same payoff symmetric state: i.e. all firms choose the same output levels. This could happen at any level of competition. Secondly, if at the start (or at anytime) there are some collusive industries, then the only possible long-run state is for all industries to be colluding. If there are some colluding industries at any time, then they will not go away. There is then no possibility that the rest of the economy can persist in a state which earns less than *JPM*.

Fig 8.6

In Figure 8.6 we depict the flows of markets between the different categories. At the top is the roach motel: the absorbing collusive state. Below that are the other above aspiration industries. On the bottom are the two mixed and below aspiration industries. Now, clearly, under the switching assumption, there are outflows from the below aspiration set to all of the others (anything is possible) every period. Likewise, there are flows from the mixed industries: again, they can go to all of the other categories, except that the flows need not be active every period (it depends exactly which pairs are involved). Lastly, there are flows from the above aspiration group to both the mixed and below aspiration groups. These flows occur not due to switching, but changes in the aspiration level: if the aspiration level rises (as average profits rise), then the current profits of firms in these industries may be below the new aspiration level. We can see that there is a continuous flow between the various categories: but every period, some firms will end up in the absorbing collusive state. There is no way for industries to escape it in the long run!

In essence, the argument is that in the long-run, all firms need to earn at least the average profit. Assuming that there are some industries (even a very small proportion) colluding at some point, then for all firms to earn average profits mean that they must collude and earn JPM.

Theorem: the Inevitability of Collusion (Dixon 2000). Suppose that at some time there are some collusive industries. Then under the assumptions, in the long all industries will collude.

This is a remarkable result. It shows that the pressure of the capital market on firms will force them to collude: competition cannot survive! Let us just think how this works, the forces involved.

First, let us think how can collusion persist, how can it be stable:? We all know the standard arguments that there is an incentive to deviate from a collusive output: one of the firms can earn higher profits if it deviates by (for example) producing a larger output. Suppose that one firm does this. Then it will increase its

own profits, but reduce the other firms profits (and reduce the combined profits). The other firm will now be below aspiration, and hence it will start to experiment: for example it may produce a larger output. Then both firms will become below aspiration and continue to experiment until both firms are above aspiration: i.e. collusive! We can think of the period of experimentation following the defection as analogous to a punishment (as discussed in *oligopoly theory made simple*). Whilst there is no sense in which the punishment is optimal, it will act in a similar fashion. The point is that although aggressive or competitive behaviour might bring a higher payoff in the short run, it cannot survive in the long run. The reason is that it will generate a response from the competitor, which will set the industry in motion until it can settle down into a situation where both firms are earning average/normal profits. This is not unrealistic: firms who are very aggressive towards rivals will become involved in price wars and similar episodes. Their shareholders might well prefer them to reap the rewards of a cosier relationship with competitors.

The implications of this result might be taken as quite far reaching: we should expect the operation of capital market pressures to enforce collusion, not competition. Competition tends to reduce profits, at least in the long run, and hence cannot be sustained in the long run. The model as presented did not include entry. However, entry *per se* need not alter the result. If there is a fixed number of firms in the industry (two or more), then the same arguments will indicate that collusion will be established between them. Now suppose that we impose an entry condition on the economy. One way to do this is to divide the model into two stages: first an entry phase and then the market phase. There is a fixed set up cost. With free entry, entry will occur to the point where expected profits are zero. If we take expected profits as the long run steady state profits (i.e. *JPM*), then the *JPM* per firm will equal the entry cost. Given the free entry equilibrium number of firms, the equilibrium will be collusive: free entry just drives the average profit to zero. It is perhaps worth looking briefly at an example of how the theorem might work out in a concrete example.

8.4.5 An example: Cournot Duopoly.

Perhaps the simplest economic application of our model is to the Cournot duopoly model with linear demand and without costs which we have considered earlier.

Recall that JPM is 0.125, and is maximised by the output pair (0.25, 0.25).

We³⁰ allowed for $K=21$ types of firm (not just the three $\{A,B,C\}$). To do this, we chose a grid of granularity 0.025 over the range³¹ 0.1 to 0.6, perturbing it slightly by moving 0.325 to 0.333 (1/3), so that the Cournot-Nash output was included. Hence $K=21$ and there are 231 possible pairs of output. We assumed that there is random switching: if a firm decides to experiment, it chooses each of the 21 strategies with a probability of 1/21.

The simulations were initiated from the initial position with a uniform distribution over all pairs. The results of the simulation are depicted in Figures 8.7a and 8.7b. In Figure 8.7b, we see the path of average profits over time: in Figure 8.7a the evolution of population proportions of the JPM market (0.125, 0.125) and the symmetric Cournot market are depicted (note that the proportions are measured on a logarithmic scale).

Fig 8.7a and 8.7b here.

From Fig 8.7b, we see that the average profits converge to the symmetric joint profit maximum of 0.125. However, the time path of profits is non-monotonic: at particular times there appear large drops in profit. The reason for this is quite intuitive. As the average profit level increases, it surpasses that of one or both firms, which start to experiment. The profits of firms at those markets will then on average fall below the population average as the firms disperse over some or all output pairs. The effect of this can be quite dramatic: the discontinuity is particularly large when a symmetric market goes critical, since both firms at each such market begin to experiment and spread across all possible output pairs. However, whilst the time-series of profits is non-monotonic and "discontinuous", there is a clear upward trend and convergence to 0.125.

From Figure 8.7a, the proportion of colluding firms ($P(S)$) is monotonic, but far from smooth. Corresponding to the discontinuous falls in population average profit, there are jumps in the proportion of firms at the *JPM* market, corresponding to the jumps in average profit. The proportion of firms at the Cournot pair (1/3,1/3) is a highly non-monotonic time series. The first thing to note is that in the initial stages of the simulation, the proportion of Cournot markets exceeds the proportion of *JPM* markets. This can occur because during this period the Cournot pair is also in the above aspiration set: until average profits reach 1/9, the Cournot pair will "absorb" markets with one or both firms below aspiration. The fact that the Cournot pair attracts more than *JPM* is due to the fact that early on more markets with experimenting firms can reach the Cournot pair than the *JPM* pair. However, after 50 iterations, the Cournot pair has a smaller proportion than the *JPM* pair, and is in the below aspiration category most of the time. The time-series of the Cournot market type is not atypical: most pairs except *JPM* have a similar time-series profile. The convergence of the proportion of markets towards type *JPM* is steady but slow: this is because the probability of hitting *JPM* from other locations is small throughout the simulation: from each market in which both firms experiment there is a probability of 1/442 of moving to *JPM*. Convergence is in general quicker with fewer strategies and non-random switching rules. We explore more specific rules (imitation, best response etc.) in the Cournot model using simulations in Dixon and Lupi (1997).

9.5 Conclusion: how economists can get smarter by making agents dumber?

In this chapter we have gone round in historical circles, traversed the surface of a donut and visited the duopoly archipelago. What have we learned? Well, I think that we can see that if we are willing to assume that economic agents are intelligent rather than having some abstract notion of "perfect rationality", we can learn quite a bit. Agents interact in a social situation and can learn, both from their own experience, and the experience of others (either their neighbours or the general population). If we assume that agents have some ethereal notion of perfect rationality, then we cannot begin to understand this. To assume that agents are perfectly rational means that we have to adopt a framework where they are able to understand what is going on in

some significant sense. However, perfectly rational agents can only solve problems we can solve, and we can solve only very simple problems. So, if we stick with perfectly rational agents, we will restrict our vision to simple models.

In this chapter I have outlined the possibility of an alternative approach. Let us assume that agents are boundedly rational: they may even be completely dumb, or just use some rules for updating that are intelligent but not in any sense optimal. In terms of *Artificial intelligence and economic theory*, we are adopting a specific model of reasoning where we specify what the agents think, how they think (I am using “think” in a broad sense here, since most economic agents are not individuals). We can then model the interaction of agents in some sort of *network*. The networks we have looked at are very simple. However, in principle we can at least imagine the economy as an extremely complicated network: a network possibly as complicated as the neural networks in the human brain or possibly even more complicated. There may have different levels of organisation: for example in the Duopoly Archipelago the capital market worked at the aggregate level, imposing the population average on the individual firms. These levels can then interact and yield interesting and novel outcomes. In the Duopoly Archipelago, if the duopolists were all playing a prisoner’s dilemma, then they would (in the long-run) be forced to collude. Thus the economy is operating in such a way that individual agents are forced to choose a strategy that is dominated. In a more general context *an agents actions in equilibrium may well be far from optimal*. Note that I am using optimal in a private sense: the individual firms are not choosing their best responses to each other (failure of private optimality) and also the outcome is not socially optimal. Whilst it is true that collusion maximizes the joint profits of the firms, the consumers (whose welfare does not appear in the payoff matrix) lose out.

In this sense I think I have answered my own challenge put forward in *Artificial intelligence and economic theory*: there may be strong forces in an economy leading agents away from optimising behaviour in strategic situations. Optimising behaviour can only survive or predominate if it earns higher profits than non-optimal behaviour. In non-strategic situations this is no longer true. As we saw in *oligopoly theory made simple*, non-profit maximizing managers may in the end earn more than profit maximizers. When ignorance is bliss, ‘tis indeed folly to be wise. *Since most*

economic interactions are indeed strategic, we should certainly not assume that agents optimise all of the time.

Economists have tended to ignore these higher level (non-local) interactions and focus on isolated pairs of players, or overlapping networks of neighbourhoods. However, in the information age the economy is becoming explicitly and consciously inter-connected: this self-knowledge imperfectly reflects and mirrors the objective interconnectedness of the economy revealed in the story of *infinity in a pencil*. Higher levels of organisation exist: in particular capital markets and to a lesser extent labour markets bring together different markets and parts of the economy. This is something that economists really need to focus on in some detail in the years ahead. Rather than pondering the deliberations of rational agents interacting alone or in isolated pairs, the focus should be more on intelligent agents interacting in social systems.

Endnotes.

¹ For the many examples of systematic and common behaviour that violates conventional axioms, see Thaler's various books: Thaler (1991, 1992, 1993)

² And other areas: for example Sargent (1993) for applications to macroeconomics.

³ Nearly all rational agents in game theory papers are female nowadays: they were mainly male before the mid-1980s.

⁴ Economists are not alone here. Philosophers have tried to do the same thing: for example deriving moral laws from some abstract notion of rationality. I doesn't work there either!

⁵ I am simplifying things rather a lot here: for a detailed and clear exposition of the process of evolution, you can do no better than reading Dawkins (1986).

⁶ The biologist Lamarck had different ideas. He thought that characteristics acquired during a parents life could be passed on. In Biology this has been shown to be incorrect. However, in terms of social learning it is almost certainly correct.

⁷ Again, whilst it is clear that Cournot was the first person to introduce the concept of the Nash equilibrium, I follow common usage in naming it after Nash. Economists used to often call the "Nash" equilibrium a "Cournot" equilibrium. With the spread of game theory in the early 1980s, this usage dropped out.

⁸ i is called the "row" player, because his choice of strategy determines which row we are on: likewise j is the column player.

⁹ This is a special case of the class of *payoff monotone* selection dynamics in which

$$g_i(t) = g \left[\frac{\Pi_i(t)}{\Pi(t)} \right] \text{ with } g' > 0.$$

¹⁰ Throughout this section, we adopt the simplification that there are only "pure" strategies and no "mixed" strategies. Game Theorists are much keener on mixed strategies than economists, the whole concept being somewhat problematic.

However, for those who want the "proper" definition, see the Weibull (1995) chapter 2.

¹¹ Note that the alien invaders are subject to the same random matching process: they do not arrive by one ship and spread out as in the film version, but arrive as random individuals.

¹² Note that if the replicator dynamics has an attractor, it need not be *ESS*: the attractor of a dynamic system might be a limit cycle or a non-*ESS* steady state.

¹³ An alternative concept of stability is *Lyapunov* stability. A steady state is Lyapunov stable if a small deviation does not lead to any further deviation (it need not actually return to the steady state as is required by the asymptotic stability concept).

¹⁴ To be accurate, we are assuming that $P = \min[0, 1 - X_i - X_j]$ to secure this result.

¹⁵ The analysis here has been brief. For a full if technical analysis of the relationships between the three concepts, see Weibull (1995, chapters 2 and 3).

¹⁶ In fact, the fancy names for all these things comes from *Graph Theory*: the houses are usually called *nodes*, and the paths connecting them *vertices*.

¹⁷ The reason it will not work is that to get the donut you need to stretch some parts and compress others. So, if you want to make a donut shape, use a sheet of stretch material. First make a cylinder and then join the ends together.

¹⁸ Well, any topologically equivalent shape to a circle, i.e. any line which does not cross itself and has no ends.

¹⁹ The two dimensionality of the earth's surface is reflected in the fact the each point on the surface can be represented by two numbers: its longitude and latitude.

²⁰ I leave out the possibilities of pollution from pesticides, fertilisers and GM crops.

²¹ The male pronoun does not reflect any presumed irrationality on the part of farmers, although it certainly helps to be crazy if you are a farmer nowadays.

²² We can see the "imitate your best neighbour" as a heuristic algorithm as discussed in *Artificial Economics and Economics*.

²³ Remember, since the top and the bottom are connected, there are two borders.

²⁴ There are obvious exceptions here, such as non-profit organisations, owner managed firms. However, all commercial organisations are covered by bankruptcy laws and the requirement to be solvent (i.e. a positive cash flow).

²⁵ Well, this needs to be adjusted for risk: the risk-adjusted rate of return needs to be equalised across all industries and firms.

²⁶ In fact, since the identity of the firms is irrelevant, we treat $\{i,j\}$ as identical to $\{j,i\}$: hence without loss of generality we write the pairs as $\{i,j\}$ with $i \leq j$.

²⁷ As in the Atkinson and Suppes (1958) "finite Markov model", where there is a probability that at time $t+1$ the firm will switch from the strategy it plays at t : the key difference with the present paper is that we use an explicit aspiration based model.

²⁸ It is straight forward to allow for firm specific aspirations.

²⁹ In fact, as the British physicist Stephen Hawkins discovered, due to weird quantum effects, black holes do radiate a bit, so matter does escape.

³⁰ I would like to thank Paolo Lupi for implementing this programme in Gauss.

³¹ We did not allow for a wider grid range (e.g. $[0,1]$), because the additional strategies are often ones with very low or zero profits: they slow down the simulation without adding any extra insight




Figure 8.1 Relationship between equilibrium concepts for steady states.




Fig 8.2. The S-curve or logistic curve.

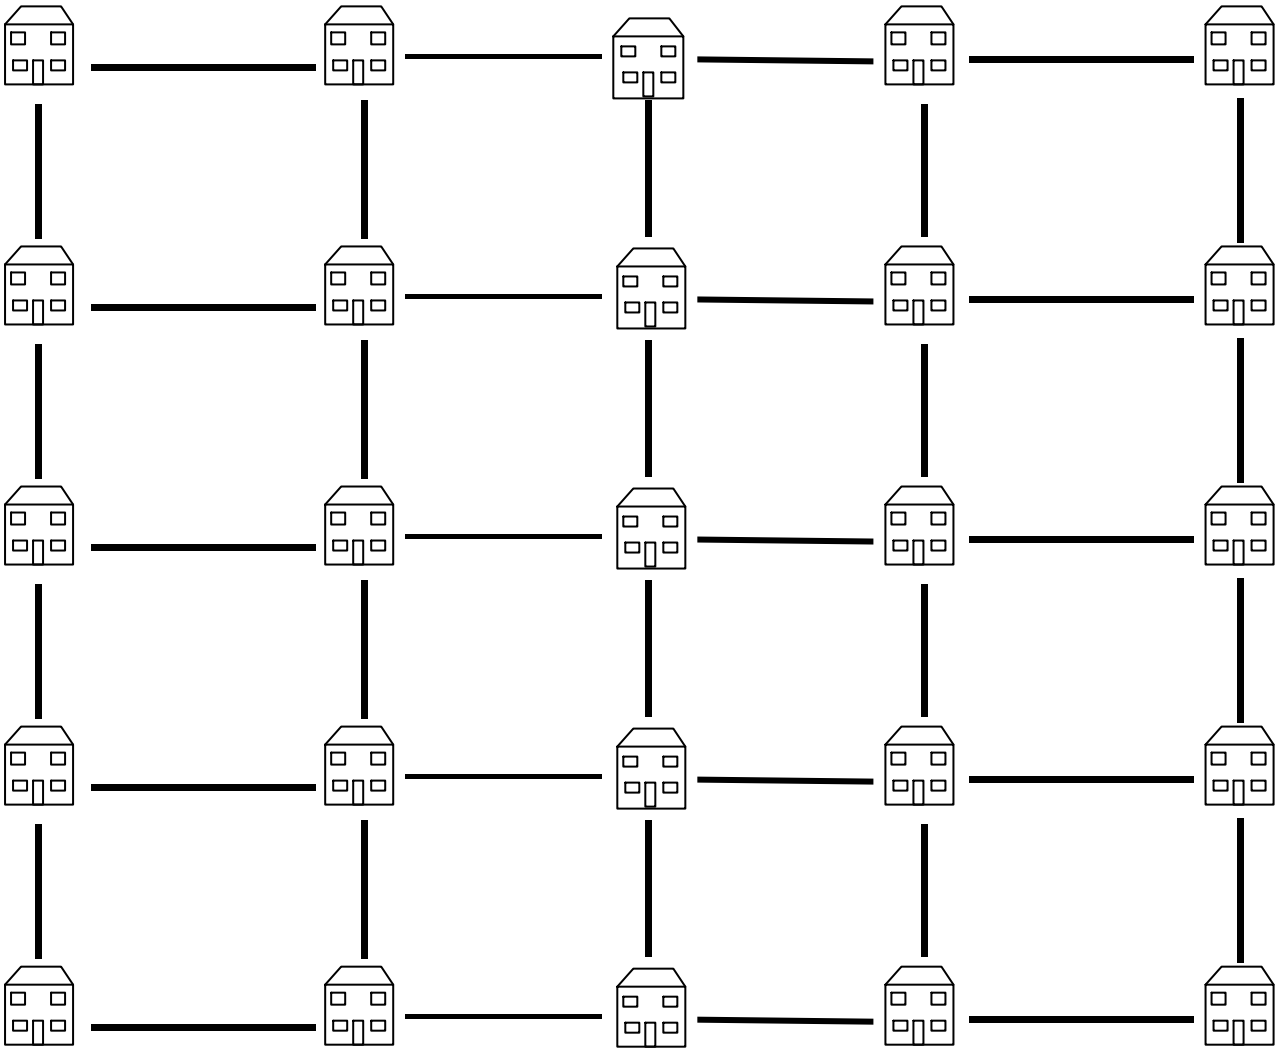


Figure 8.3. Donut World




Figure 8.4. Diffusion in Donut World

D	D	D	D	D	D	D
D	D	D	D	D	D	D
D	D	D	D	D	D	D
D	C	D	C	D	C	D
C	C	C	C	C	C	C
C	C	C	C	C	C	C
C	C	C	C	C	C	C

Figure 8.5. Cooperation and Defection in Donutworld.

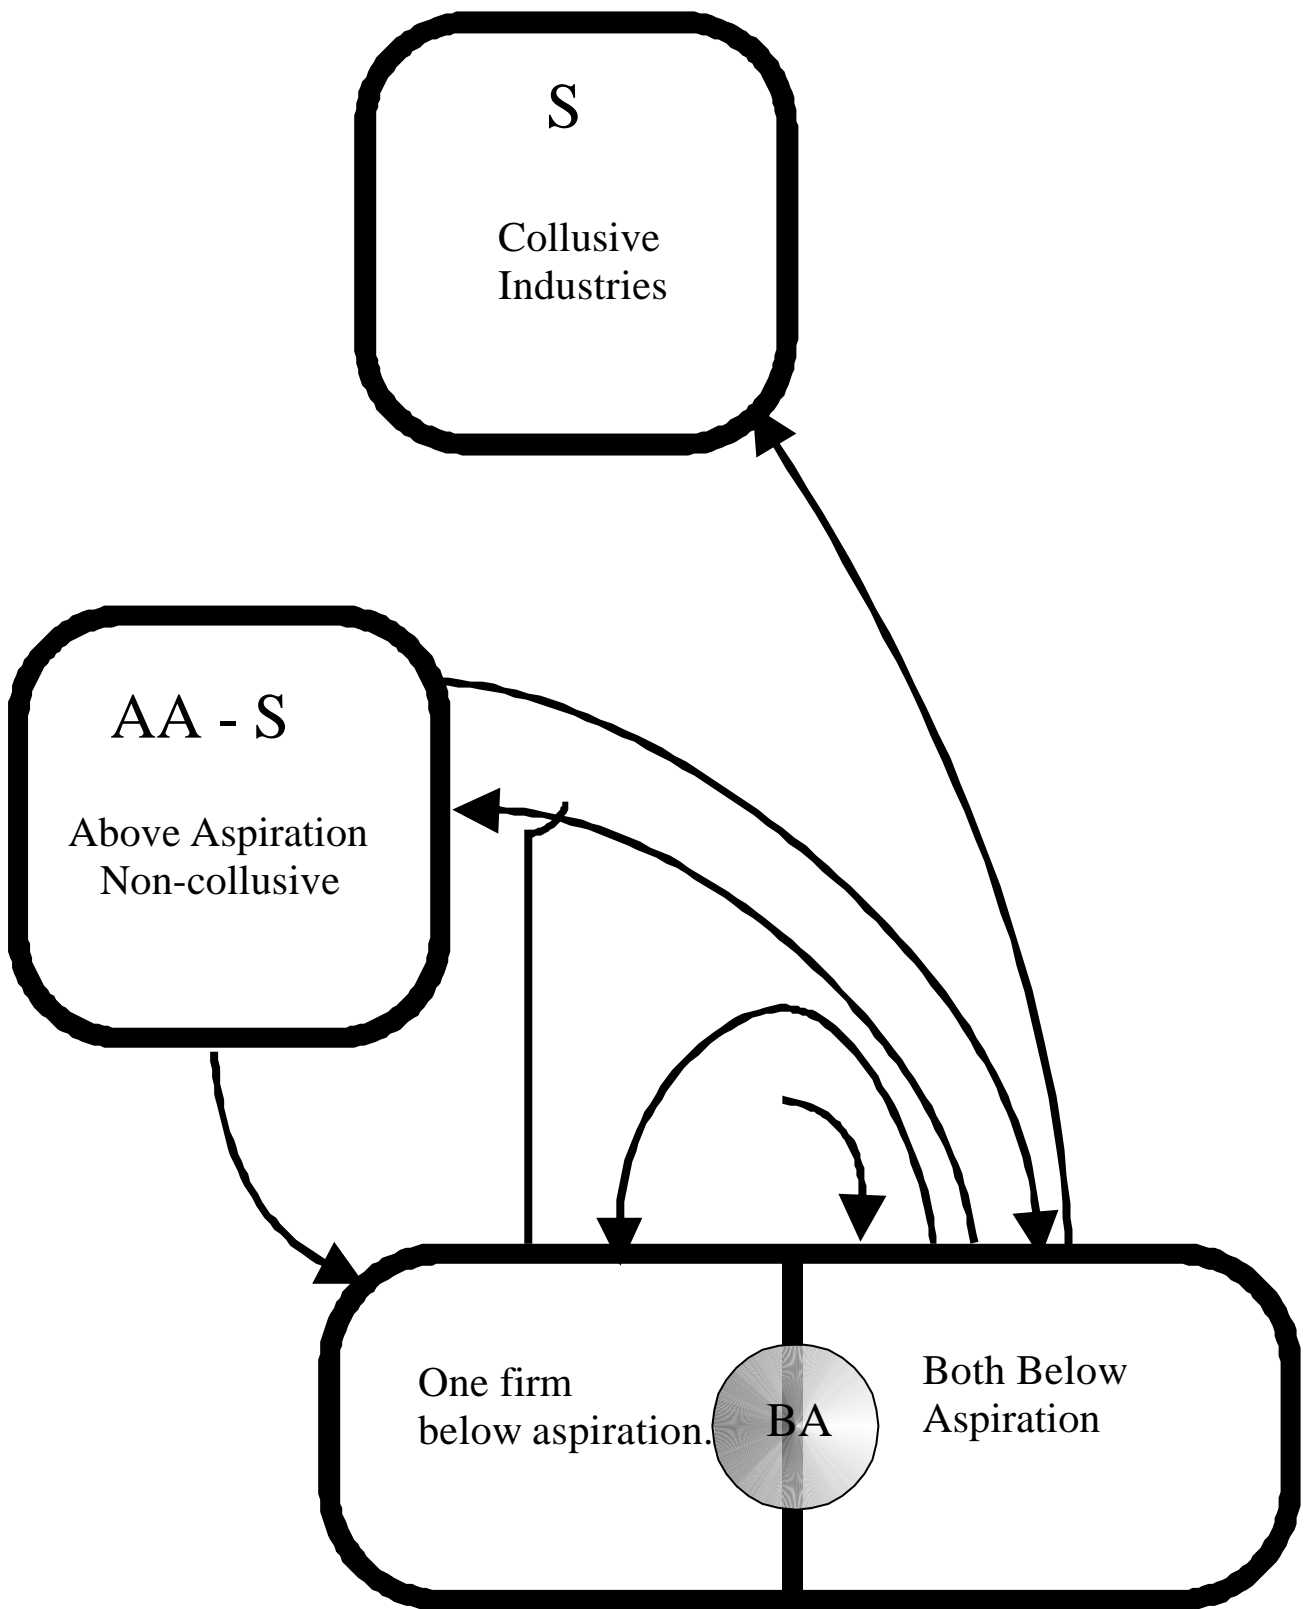


Figure 9.6 Flows of Industries between aspiration states.