

Chapter 7: Some Thoughts on Economic Theory and Artificial Intelligence

7.1 Introduction.

In this chapter I will offer some thoughts on the potential contribution of Artificial Intelligence (AI) to economic theory. I write as someone who is at present an orthodox economic theorist and who has only recently been introduced to the ideas and achievements of modern artificial intelligence. My overall impression is that artificial intelligence does have a potential to offer economic theory, both in terms of providing a fruitful perspective from which to view currently problematic issues, and through raising new and interesting problems that have been hitherto marginalized or ignored.

In my view, the main potential contribution of artificial intelligence to economic theory lies in providing a practical methodology for modelling reasoning, and hence rationality. Economic decision-making is one of many human activities which can be said to display ‘intelligence’, since it involves potentially complex reasoning and problem-solving. Artificial intelligence is a branch of computing which aims to design machines which can perform such ‘intelligent’ activities; as such, it has much to say about how humans go about solving problems. There now exists a wide body of research on artificial intelligence in a variety of applications which may well prove useful to economic theorists.

However, the current orthodox approach of economic theory to rationality is well-established and simple. Advocates of artificial intelligence need to demonstrate clearly why there is a need to change perspective, and how it will pay off in terms of theoretical understanding. Therefore in Section 7.2 I explore both the current orthodox model of ‘rationality without reasoning’, and in particular I consider to what extent it can be extended to embrace ‘bounded rationality’. As I will argue, I believe that by generalizing strict optimization to approximate optimization we can maintain the simplicity of the current orthodoxy without needing explicitly to model problem-solving with artificial-intelligence techniques. In section 7.3 I take the more positive approach of considering how artificial-intelligence techniques can make a useful contribution in modelling the behaviour of economic agents in complex decision situations. In particular, I describe the use of finite automata to model the complexity of strategies, and also the issues arising from the choice of strategies. To put matters

very simply, insofar as agents tend to ‘get things right’ in the sense of choosing an action or solution close to the optimum, we (as economists) need not really worry about how the decision is arrived at. The role for artificial intelligence seems to me to have the greatest potential in situations where agents make mistakes. This potential is, of course, of great relevance to the study of disequilibrium, an area of economics that has proven notoriously difficult within the orthodox model of rationality. Lastly, in Section 7.4, I consider the implications of modelling reasoning in strategic environments. Insofar as the method of solving a problem might influence the action you choose, reasoning itself can be a method of precommitment. This is illustrated in the context of Cournot duopoly.

I would like to acknowledge my debt to the writings of Herbert Simon: whilst I have not referenced his works explicitly or in detail, his ideas were formative in my education and have clearly influenced the way I look at things.

7.2 Orthodox Economic Rationality: Rationality without Reasoning

At the heart of orthodox economics beats a model of rationality. The economic conception of rationality has become embodied by the mathematical theory of constrained optimization. Put at its simplest, a rational agent is conceived of as choosing the best option open to him, given his constraints. The historical basis of this lies in utilitarianism, and in particular the Utilitarian psychology of individual action, which saw mankind placed ‘under the governance of two sovereign masters, pain and pleasure’ (Bentham, (1789, p. 33)). Although more recently economics has eschewed utilitarian psychology due to problems of measurement of cardinal utilities, the notion of maximization itself has become, if anything, more central. If we consider the salient features of economic rationality, three points are worth highlighting in the present context.

(1) There is no modelling of reasoning. The process or procedure of reasoning is viewed as unimportant in itself, and the outcome inevitable. If a solution (maximum) exists, then the rational agent will arrive at that solution; agents costlessly choose their best option.

(2) Technical assumptions are made to ensure that a well-defined maximum exists. This usually involves two sorts of assumptions: first, to ensure that a well-defined continuous objective function can be specified; second, to ensure that the agents’

choice set is compact¹. It is often easy to forget how many assumptions are made by economists just to ensure both of these criteria are satisfied, so as to ensure the existence of a maximum.

(3) If it is not possible to impose plausible restrictions to ensure the existence of well-defined maximum, there exists no obvious, typical, or generally accepted solution to the problem.

Let us illustrate this with reference to standard consumer theory. A household is assumed to have preferences over possible outcomes, which are taken to be suitably defined bundles of commodities which it might consume. These preferences are 'represented' by a utility function, which assigns to each bundle a real number, with the property that bundles which are preferred have higher numbers. If we turn to the first issue, a well-defined utility function is ensured by four assumptions: (i) the 'sanity clause' of reflexivity that each bundle is at least as good as itself; (ii) that preferences are complete, so that any two bundles can be compared to each other; (iii) that preferences are transitive, so that if A is preferred to B, and B to C, then A is preferred to C; (iv) continuity, so that if one bundle is strictly preferred to another, there, there is a bundle in between that is also strictly preferred. It should be noted that the assumptions of transitivity and continuity in particular are required only for mathematical reasons, rather than any fundamental notions of rationality. The second issue of compactness is rather easier to ensure: the household is restricted to choosing its utility-maximizing bundle from its budget set, which is closed and bounded (at least if all prices are strictly positive).

The main feature of the orthodox model of economic rationality is that there is no modelling of the reasoning processes of agents, or of how decisions are reached. The formal maximization process is simply stated, and if a maximum exists it is assumed that an action yielding the maximum is chosen. This can be seen as a model of rationality without reasoning, a 'black box' or even 'empty box' approach to behaviour. The great attraction of the approach is that it means that economic theory can largely ignore the potentially complex and diverse processes of reasoning and decision-making underlying the behaviour of individual economic agents and organizations. This yields a very simple model of rationality. Whilst it has been criticized for its very simplicity, in that it ignores much real-world complexity, it has certainly delivered results and proven itself in many contexts.

In order to convince economists of its value, advocates of artificial intelligence need to demonstrate that there is both a clear need and payoff to model reasoning itself.

Whilst simplicity is perhaps the most cogent defence of orthodox economic rationality, another important defence is that it can easily be extended to encompass bounded rationality. Given that some economic environments are complex and uncertain, so that strictly defined optimizing may be inappropriate, it still may be unnecessary to deviate from the orthodox model. Whatever line of reasoning, search procedure, or algorithm agents may use, rational agents may be reasonably assumed to get 'near' to the strictly defined optimum. A very simple way to model this is to assume that agents ϵ -optimize: they choose an action that yields them within ϵ of their best payoff. Let us specify this a little more formally. Suppose that an agent chooses an action, a , that is chosen from the closed interval $[0, A]$, and that its payoff function $U_i: [0, A] \rightarrow \mathbb{R}$ is continuous. Then there exists a maximum, U^* , which is yielded by some a^* :

$$a^* = \operatorname{argmax} U(a)$$

This is depicted in Fig. 7.1, where the optimal a^* is assumed unique. It might be argued that in a particular context the agent will not be willing or able to inevitably choose a^* .

Figure 7.1 e-Optimization

Rather, the agent might adopt some heuristic search procedure or rule of thumb that might get reasonably close to optimum in terms of payoff. The details of the precise line of reasoning taken need not (it can be argued) concern the economist. Rather, whatever the route taken, the agent can be said to be ϵ -maximizing, or choosing actions which yield payoffs within ϵ of the maximum. Given the continuity of $U(\cdot)$, there will in general be a set of such acceptable solutions, defined for $\epsilon > 0$ by:

$$A^\epsilon = \{a \in [0, A] : U(a) \geq U(a^*) - \epsilon\}$$

This set is depicted in Fig. 7.1. Two points need making. Firstly, the concept of ϵ -maximization is a generalization of strict optimization: a^* is always an acceptable

solution, and when $\epsilon = 0$ it is the only acceptable solution. Secondly, the choice of ϵ is more or less arbitrary, although it is usually taken to be more acceptable if it is ‘small’.

A simple generalization of the orthodox approach is thus sufficient to capture some elements of bounded rationality without the need explicitly to model reasoning. As in the case of strict optimisation, approximate optimization goes straight from the statement of the problem to the set of approximate solutions. Corresponding to the notion of ϵ -maximization is the notion of an ϵ -equilibrium. For example, a Nash ϵ -equilibrium is defined in much the same way as a ‘strict’ Nash equilibrium. Suppose now that there are two agents, $i = 1, 2$, choosing actions a_i from compact strategy sets A_i , and continuous payoff functions $U_i(a_1, a_2)$. An ϵ -equilibrium occurs when both players choose actions (a_1^*, a_2^*) which yield them within ϵ of their best payoff given the action of the other player. Formally, for player 1 (and analogously for 2) (a_1^*, a_2^*) satisfies:

$$U_1(a_1^*, a_2^*) \geq U_1(a_1, a_2^*) - \epsilon \quad \text{for all } a_1 \in A_1$$

This can be represented by “reaction correspondences”, where a reaction *correspondence*² gives the player actions that yield within ϵ of his best response given the other player’s action. For player 1, we have $r_1: A_2 \Rightarrow A_1$:

$$r_1(a_2) = \{a_1 \in A_1 : U_1(a_1, a_2) \geq \max_{a_1} U_1(a_1, a_2)\}$$

These reaction correspondences are ‘fat’ because there is a set of ϵ -maximal responses to any action by the other agent. An ϵ -equilibrium occurs at any point in the intersection of the two reaction correspondences, as depicted in Fig. 7.2. It should be noted en passant that in general there will be many (in the ‘continuum’ sense) ϵ -equilibria, and they may well be Pareto-ranked. To see this, merely consider what happens as ϵ gets large: for ϵ large enough, any conceivable outcome will be an ϵ -equilibrium! However, whilst multiplicity is endemic to ϵ -equilibria, it is not exclusive. There is no real reason why there should be unique strict optima and unique equilibria in strict Nash equilibria. They are usually ruled out for convenience sake, because models with unique and well-behaved equilibria are simpler to deal with.

The concept of bounded rationality as embodied in the notions of approximate optimization and equilibria, whilst not in general usage, is certainly not uncommon. Well known recent examples include Akerlof and Yellen's work on business cycles (Akerlof and Yellen (1985)), and Radner's work on cooperation in the finitely repeated prisoner's dilemma (Radner (1980)). The notion of e-equilibrium has been often used when no strict equilibrium exists (see, inter alia, Hart (1979), Dixon (1987)).

The notion of e-optimization is a simple generalization of strict-optimization that maintains its parsimony. The details of reasoning need not be considered, and it is not mathematically difficult to deal with. However, there still seems to me to be further conceptually simple extensions of the more-or-less orthodox that can be used to try and capture aspects of bounded rationality without the need to model reasoning. One possibility that has not (I believe) been explored, is to treat the rational player's decision as a type of mixed strategy. The decision-making process might have general properties: it would be more likely to choose a better outcome; it might be able to avoid payoffs that are particularly bad, and so on. Again, this would be a generalization of strict-optimization (which puts a probability of one on the maximal, and zero elsewhere).

Figure 7.2 Nash e-equilibria

Whilst this approach has not been pursued, I believe it to be indicative of the fact that there are many possible extensions to the orthodox model that maintain its key advantages of simplicity and parsimony. The task for artificial intelligence is to show that it can yield something more than can be obtained by developing the orthodox approach.

7.3 Complexity and artificial intelligence

I have described orthodox economic rationality as 'rationality without reasoning'. If a problem is sufficiently simple and well defined so that there exists a solution which is easily computable, the precise method of solution adopted may not matter. For example, a non-singular square matrix can be inverted by different techniques, all of which will yield the same solution if applied correctly. However, for complicated problems it may be the case that although we know a solution (optimum) exists, we do not know how to find it with certainty. Even worse, we may not know if a solution

exists at all. In such a situation, the precise form of reasoning may be important, because it will determine the types of actions or decisions arrived at. In this context I am equating ‘reasoning’ with a particular method of searching for a solution (an algorithm). More importantly, different methods of searching for a solution may tend to yield different types of outcomes. In this context the choice of reasoning itself can become a strategic decision, as I will discuss in the next section.

There are two dimensions of complexity that deserve the particular consideration of economists, which I will outline in this section. Firstly, there is the complexity of *strategies*; secondly, there is the complexity of *choice of strategies*. I shall now briefly discuss these two issues.

Taking a cue from standard game theory, a strategy can be seen as a rule for choosing actions, in which the action chosen is a function of the information available to an agent. At any time, the ‘information’ available might be taken to consist of all of the past actions of various players, plus the realizations of various exogenous parameters. For example, an oligopolist at time t might know the history of play up to time t in the sense of knowing all the choices of outputs by itself and other firms, as well as (for example) the past realizations of demand. If we denote this ‘history’ at time t as h_t , a strategy is a rule (or in mathematical terms, a mapping) that tells the firm at any time t what action to take given h_t :

$$x_t = r(h_t, t)$$

In informal terms, the strategy can be seen as a ‘game plan’ for the firm, telling it what to do in every possible contingency. In standard game theory, there is no limit to the complexity of the strategies chosen by agents. Furthermore, some of the results of game theory require very complicated strategies to support equilibria (for example, the ‘carrot and stick’ punishment strategies supporting equilibria in Folk theorems).

Various authors (Rubenstein (1986), Abreu and Rubenstein (1988), Kalai and Stanford (1988)) have argued that some limits or constraints need to be put on the complexity of strategies chosen. A popular way of capturing ‘complexity’ in this context is to use the notion of a ‘Moore machine’ or ‘finite automaton’ as a way of representing a strategy. A finite automaton is a machine with a finite set of states, and a specified initial state. It then has two rules: one specifies what action it takes at time t as a function of its state at time t ; the second is a transition rule which specifies what

its state in time $t + 1$ will be as a function of both its state in time t and the action of the other player at time t . The size of the automaton (its ‘computing power’) is captured by the number of states it has: ‘... the complexity of a strategy will be identical with the size (number of states) of the smallest automaton which is able to implement it’. (Kalai and Stanford 1988).

This can easily be illustrated by a couple of examples using ‘transition diagrams’ to represent the automaton (see Rubenstein (1986) for a full and lucid exposition). Let us take the standard prisoner’s dilemma game, where there are two agents, $i = 1, 2$, each with two strategies, C (cooperate) and D (defect), with payoffs as in Table 1. The automaton can be represented by $\langle Q_i, q^0, ?_i, \mu_i \rangle$ where Q_i is the set of states; $q^0 \in Q_i$ is the initial state; $?_i$ gives the action a_i as a function of state q_i , and μ_i is the transition rule giving the next period’s state. The simplest strategy is generated by a ‘one-state’ machine, where $Q_i = q^0 = Q$, $\mu(Q) = Q$, and either $?_i(Q) = C$ or $?_i(Q) = D$. That is, a ‘one-state’ machine can only cooperate (or defect) all of the time: because it only has one state it can only ever choose one action, as depicted in Fig. 7.3a. In order to respond, the automaton requires more states and more complicated transition rules. Let us consider a two-state machine, with one state, q^D , for defection, and one, q^C , for cooperation. The ‘tit-for-tat’ strategy is when the player cooperates until the other player defects, in which case he punishes the defection by playing D himself for one period before returning to C (see Axelrod (1984)). The automaton implementing this strategy is represented in Fig. 7.3b, and is defined by a set of states $Q = \{q^C, q^D\}$, initial state $q^0 = q^C$, action rule $?(q^a) = a(a = D, S)$, and transition rule $\mu(q, a) = q_a(a = D, S)$. The ‘grim’ punishment strategy punishes a defection forever, and is represented in Fig. 7.3c. The ‘grim punishment’ automaton has an initial state, q^C . If the other automaton ever plays D, the machine switches to its defect state q^D , where it remains forever. The main point is that the automaton with more states can implement more complicated strategies.

Figure 7.3 Strategies as finite automata

		Player 2	
		C	D
Player 1	C	2,2	0,3
	D	3,0	1,1

The Prisoner's dilemma (PD)

Using the model of a finite automaton to represent a strategy, Rubenstein et al have represented a game as involving agents choosing machines rather than strategies. The players 'optimize' over the choice of their machines (strategies), and then the machines play the game. As such, this approach ignores issues of computational complexity, and focuses merely on the issue of implementational complexity. Abreu and Rubenstein (1988) draw the analogy of sophisticated managers formulating simple operating rules for the firm.

The issue of computational complexity is central to artificial intelligence. Some problems are simple enough to have a deterministic method which will (nearly) always arrive at a solution. Consider the 'travelling salesman' problem:

A salesman has a list of cities, each of which he must visit exactly once.

There are direct roads between each pair of cities on the list. Find the route the salesman should follow so that he travels the shortest possible distance on a round trip, starting at any one of the cities and then returning there.

The solution to this problem can be found simply by exploring the 'tree' of all the possible paths, and picking the shortest one. This 'brute force' method will work, but it becomes very expensive in terms of computational requirements. With N cities, there are $(N-1)!$ different routes: each route has N stretches. The total time required to explore the entire tree is then of order $N!$ – with only 10 cities this method clearly becomes very lengthy ($10!$ is 3,625,800). This is known as the phenomenon of *combinatorial explosion*. Since computations are not costless, rational agents need to trade-off the computational cost of a decision or search procedure with the benefits in

terms of the eventual payoff. The notion of efficiency is important here: a search process is more efficient if it obtains a higher payoff (on average) for the same or fewer computations (on average). Evolving efficient search strategy involves using *heuristics*, which are guides, short-cuts or rules of thumb that we believe will get us near enough to the solution with reasonable computational requirements. For example, the travelling salesman problem can be solved using the useful general purpose heuristic of the *nearest neighbour algorithm*:

1. Select any city as your starting point.
2. To select the next city, consider the cities not yet visited. Go to the city closest to the one you are currently at.
3. Repeat 2 until all of the cities have been visited.

This algorithm requires far fewer computations than ‘brute force’: we need to visit N cities, and at each stage we need to consider the distances between where we are and the as yet unvisited cities, of which there will be on average $(N-1)/2$. Computational time is therefore proportional to $N(N-1)/2$, or more simply N^2 , which is far superior to $N!$ The nearest neighbour algorithm need not yield the shortest route: Bentley and Saxe (1980) have found empirical evidence that when cities are distributed at random it performs on average about 20% below the optimum. Much of practical artificial intelligence involves the construction and evaluation of heuristic search procedures (see, for example, Polya’s classic (1957)).

One response to the issue of computational cost and decision-making is to maintain the notion of maximizing behaviour, but simply to add in a constraint reflecting costs of decision, computation or whatever. This may be reasonable in particular applications. For example, much of standard theory assumes that it is costless to change or adjust variables. It is simple enough in principle to introduce costs of decision or adjustment into our economic models (see, for example, Dixon (1987) for a discussion of menu costs). However, this approach cannot answer the general problem. If one is unable to solve a simple optimization problem X , it is unlikely that one can solve the more difficult problem, of optimizing X subject to some computational constraint. This sort of *Super-optimization* is not the answer to an initial failure of a simpler optimisation problem. Rather, agents are forced to adopt reasonable decision procedures rather than the optimal.

The notion that many fundamental economic decisions are complex and uncertain has of course a long pedigree in economics, with its notions of rules of thumb (see Cyert and March (1963), Hall and Hitch (1951), Simon (1957) *inter alia*). However, as I argued in the previous section, we can relax the assumption of strict-optimization to e-optimization, which still largely maintains the ‘empty box’ methodology of rationality without reasoning. If agents use various heuristic techniques or rules of thumb, then presumably they will yield payoffs that are close to the optimal in some sense. In this case the precise nature of the heuristic need not bother us. The argument is really no different than in the case of ‘perfect rationality’, where we can predict the optimal choice irrespective of the method of solution. The only difference is that whereas the orthodox approach predicts that agents arrive at an optimal solution, we can relax this to a prediction that the agent will arrive close to the solution.

If economics is to abandon its model rationality without reasoning, it needs to be shown that there is a need to look at reasoning itself. In a complex decision problem, we may not be able to find a solution or optimum with certainty, and indeed may not even know if a solution exists. The method of reasoning, or searching for a solution may in these circumstances be important, because it will determine the actions and decisions of agents, and hence different methods may yield (or tend to yield) different types of outcomes. If we are not confident that the method of reasoning will tend to yield solutions close to the optimum, then the matter is different. I believe that the method of reasoning becomes most important when we need to understand and explain why agents make decisions that deviate considerably from the optimum. Almost paradoxically, reasoning is important only when it leads to mistakes. We only need to understand the mechanics of the rational agent’s decision process when they fail.

Let me illustrate this with the example of Section 7.2, where an agent has to maximize a continuous function, $U(a)$, over the interval $[0,A]$. Suppose that the function is as depicted in Fig. 7.4. We can see that analytically there exists a unique global optimum at a^* : this is the choice predicted by standard economic theory. Approximate optimization would perhaps predict a point close to a^* . However, suppose that our agent is in the situation of having to optimize without knowing what U looks like. He can compute the particular value of U at particular point $a \in [0A]$, and its gradient (or whatever) if defined, but only at a cost. The problem is rather like

that of an econometrician trying to maximize a complicated non-linear likelihood function. There are different ways of going about this sort of problem, all of which are considered 'reasonable'. Several methods are variants of hill-climbing algorithms, as used from a point a_0 chosen at random (or by informed guess). You then compute both the value of the function and the gradient at that point: $U(a_0)$ and $U'(a_0)$. You then move a certain distance (perhaps specified by you) in the direction of steepest ascent. You stop when the function is sufficiently 'flat' and concave: usually this is defined by some predefined tolerance $d > 0$, so that 'flat' means $|U'| < d$. Depending on the costs of computing relative to likely gains, you may wish to start several ascents from different points. Two points are worth noting about such search procedures. Firstly, they will almost always fail to reach the solution A^* : $\{a^*\}$ is a singleton in the interval $[0, A]$, and is of measured zero, and hence in a loose but clear way, a^* will almost certainly never be chosen. However, as more and more points are computed, the sample maximum will tend towards the global maximum (this is ensured by the continuity of $U(a)$). For a survey of econometric applications see Quandt (1983, chapter 12).

Figure 7.4 Explaining a mistake

The shortcomings of hill-climbing algorithms are well known (and concern 'spikes', 'ridges' and 'plateaux'). It is clearly an 'intelligent' search process that is more efficient than random search. However, depending on what the functions to be maximized look like, hill-climbing may or may not be expected to get close to the optimum. Let us consider the example of Fig. x: there are three local optima $\{0, a^*, a^{**}\}$. If the agent starts to hill-climb in the interval $[0, a^1]$ he will tend towards 0; in the interval $[a^1, a'']$, he will tend towards the global optimum a^* ; if $[a'', A]$ he will tend towards a^{**} (assuming that at points $\{a^1, a''\}$ the hill-climbing routine is equally likely to go in either direction). Our prediction of the agents eventual choice of action would depend upon the number of computations available. However, if $[a^1, a'']$ is small relative to $[0, A]$, we would certainly need to put a positive probability on all actions close to each local optimum. Furthermore, as drawn, it is clear that if only a few computations are made, then it is much more likely that the largest value computed will be close to $U(a^{**})$, since $[a'', A]$ is much larger than $[0, a^1]$ or $[a^1, a'']$.

Suppose that we observed an agent choosing action a^{**} , how might we explain it? Orthodox strict-optimization would be powerless: the optimum is $U(a^*)$, and it has not been chosen. The mistake is inexplicable. In practice, no doubt, the route of super-optimization would be pursued: the agent had a set of priors over $U(\cdot)$ and chose a^{**} to maximize the expected payoff. However, to repeat: *super-optimization is not an adequate response to the failure of optimization in the face of computational complexity. If you cannot solve a simple problem, it is unlikely that you can solve a more difficult one!* However, if we abandon the option of rationality without reasoning, matters are easier to explain: 'our agent adopted a hill-climbing algorithm. Given a limited number of computations this was quite likely to end near a^{**} . This can be explained even though a^{**} is nowhere near the optimal choice a^* , and $U(a^{**})$ is only half $U(a^*)$.

It is worth pausing here to state this argument so far, and put it in context. First, economists need not concern themselves with how economic agents solve problems if those agents successfully optimize or near optimize. We can explain and predict their behaviour as the solution of an optimization or e-optimization problem. If, however, agents make 'mistakes' by choosing actions that are far from the optimal, and/or yield payoffs significantly below the maximum, matters are rather different. Then, in order to explain the specific choice, we will need to model the reasoning underlying the choice. I have given the example of a hill-climbing algorithm yielding a sub-optimal local maximum. Again, if we consider applying the nearest neighbour algorithm to the travelling salesman problem, from some starting points it will yield terrible solutions. The role for artificial intelligence in economics would then seem primarily to be in situations where economic agents make mistakes, and possibly bad mistakes. This is in some ways a paradoxical role for artificial intelligence.

However, it is a role with great potential, not least in modelling disequilibrium. I have discussed the concept of equilibrium elsewhere in *Equilibrium and Explanation*. There are perhaps three properties that define equilibrium: firstly agents' actions are consistent (in some sense the actions of different agents 'add up'); secondly, agents are behaving optimally in equilibrium, and so have no incentive to deviate from their equilibrium actions; and thirdly, the equilibrium is the outcome of some adjustment process. If we focus on the second property, in a Nash equilibrium, each agent's actions are optimal given the actions of other agents. In disequilibrium, however, agents' actions need neither be consistent, nor optimal. This causes agents to revise and

adjust their behaviour, which may (or may not) drive the economic system under consideration towards equilibrium. It is the essence of disequilibrium that agents make mistakes. For this reason, the analysis of disequilibrium has been very problematic for economics. There seems to me to be a role for artificial intelligence in modelling disequilibrium systems, by specifying the decision rules used by economic agents. The firm, for example, can be viewed as an 'expert system' which will have some capacity for performing well in a variety of equilibrium and disequilibrium situations, but which may perform badly in others. Indeed, the standard 'myopic' adjustment rule used by Cournot in his analysis of stability can be thought of as just such a decision rule. The firm treats the output of the other as fixed and optimizes against it. In disequilibrium this may not be a good decision rule, although in equilibrium it may be 'reasonable'.

7.4 Reasoning as Precommitment: An Example

In the previous section I argued that artificial intelligence has a role in economics to explain how agents make mistakes in disequilibrium. In disequilibrium a perfectly reasonable decision rule may lead an agent to make sub-optimal decisions. As agents adjust their behaviour in response to such mistakes, there will (perhaps) be a movement towards equilibrium. In this section we will reverse the line of reasoning. In a strategic situation (e.g. oligopoly), there may be an incentive for firms to make 'mistakes'. In this case, agents may wish to adopt forms of reasoning that lead to actions which are in some strategic sense 'desirable', although they might in another sense not be optimal.

Perhaps the most important impact of artificial intelligence on economics will be that in modelling reasoning, it brings reasoning itself into the domain of choice, and hence opens it to strategic considerations. If an agent is perfectly rational, his behaviour is in a sense thereby restricted to a particular action (or set of actions), and hence becomes predictable. Given that a firm's objective is to maximize profits, it will choose its 'optimal' profit-maximizing price/output. Even if it were in the strategic interests of the firm to do otherwise, the rational firm is 'unable' to do anything other than the optimal. This is essentially the insight that lies behind the concepts of subgame perfection and dynamic inconsistency. In each period, agents

are restricted to behaving optimally; this fact can then be used to predict their behaviour and hence the future course and outcome of play.

However, suppose that we drop the assumption of rational intuition, that if a solution exists to a problem the rational agent intuitively finds it directly. Suppose instead that an agent has to choose how to solve a problem. The choice of how he chooses to solve a problem, his decision rule, will determine (to some extent) his eventual choice of action. Economic agents can therefore use their choice of decision algorithm as a form of precommitment to certain actions. As is well known, in a wide class of games there is an incentive for firms to precommit themselves.

This is perhaps best illustrated by an example, for which I will use Cournot duopoly. I have discussed this elsewhere, in terms of oligopoly (*Oligopoly theory made simple*) as well as its general significance as an equilibrium concept (*Equilibrium and explanation*). There are two firms, $i = 1, 2$, who choose outputs $X_i = 0$. Given these quantities, the price P clears the market via the inverse demand curve $P(X_1 + X_2)$, giving each firm i 's profits as a function of both outputs (assuming costless production):

$$U_i(X_1, X_2) = X_i P(X_1 + X_2)$$

A Nash equilibrium is defined as a pair of outputs (X_1^*, X_2^*) such that each firm is choosing its profit-maximizing output given the other firm's output. Formally:

$$X_1^* = \operatorname{argmax} U_1(X_1, X_2^*)$$

and similarly for firm 2. In this sense then, neither firm has an incentive to deviate given the other firm's choice. This is often represented in terms of reaction functions. Firm 1's reaction function, r_1 , gives its profit-maximizing output as a function of firm 2's reaction function (and likewise for firm 2):

$$X_1 = r_1(X_2) = \operatorname{argmax} U_1(X_1, X_2)$$

The Nash equilibrium (X_1^*, X_2^*) occurs where both firms are on their reaction-functions, i.e. $X_1^* = r_1(X_2^*)$ and $X_2^* = r_2(X_1^*)$. This is depicted in Fig. 7.5 at point N. Without precommitment, both firms have to be on their reactions-functions, since it is assumed that firms are rational optimizers. However, if a firm can precommit itself to

take any action, then it need not be on its best-response function. As is well known, if firm 1 can precommit itself to a larger output than X_1^* , it can increase its profits by moving down the other firm's reaction-function. Under standard assumptions the maximum profit for firm 1 to earn is at its Stackelburg point, S, to the right of N. At S, firm 1 is earning higher profits than it earned at N. There is thus an incentive to precommit. However, in the absence of precommitment, X_1^S is not a credible output for firm 1 to produce, since X_1^S is not the profit-maximizing response to X_2^S (which is X_1^1). In the absence of some form of precommitment, both firms are 'restricted' to being on their reaction-functions, which result in the Nash equilibrium, N.

Figure 7.5 Cournot duopoly

In standard economic models, with perfectly rational agents, precommitment has tended to be thought of in terms of some irreversible act or expenditure (e.g. investment in Brander and Spencer (1983), or delegation in Vickers (1985)). However, in the case of bounded rationality, matters are rather different. *Firms can choose decision-making rules that tend to yield certain outcomes.* For example, in Cournot duopoly the firms have an incentive to precommit to an output larger than their profit-maximizing Nash output, since this moves them towards their Stackelburg point. Thus firms might wish to adopt decision algorithms that tend to yield large outputs, that result in systematic over-production relative to the 'optimum'. For example, if firms adopt some sort of hill-climbing algorithm, they can bias the solution to be above the optimum by tending to choose large outputs as initial positions. Such algorithms need to be told not only where to start, but when to stop. As mentioned in the previous section, the latter can be specified in terms of a threshold gradient: stop searching when the gradient falls below a certain level, $|U_i'| < d$. By starting from relatively large outputs and choosing a large d , the firm can precommit itself to choosing relatively large outputs.

Conclusion

In this paper I have sought to achieve two objectives. Firstly, to state and defend the orthodox model of economic rationality. In particular, I wanted to explore the extent to which the orthodox approach has been and can be extended to embrace the notion of bounded rationality. Secondly, given this extended notion of orthodox rationality, I

sought to explore what role artificial intelligence might have in economic theory. To conclude, I will simply summarize and restate the arguments of the paper in a schematic form.

Orthodox economic rationality is a model of rationality without reasoning. Insofar as economic agents tend to get things right – or almost right – we do not as theorists need to model how they solve their constrained-optimization problems. In most economic models it is assumed that agents are ‘strict’ optimizers, who effortlessly optimize. Whilst this is an idealization/simplification, it can easily be generalized to embrace bounded rationality by adopting the notion of e-optimization. In neither case is it necessary to consider in detail how agents actually decide what to do, their ‘reasoning’. This is an advantage insofar as it means that economic theorists can avoid the complexities of the psychological and bureaucratic decision processes within individuals and organizations, and simply consider the objective problems (and solutions) themselves.

Given this extended notion of orthodox rationality, what role is there left for artificial intelligence? If orthodox rationality can handle decisions that yield optimal or near-optimal outcomes, it would appear that the main area for artificial intelligence to make a distinctive contribution is in situations where agents do not take decisions that yield optimal or near-optimal outcomes. I have highlighted two particular areas of possible research where this may be necessary: disequilibrium and strategic environments. In disequilibrium environments it is of the essence that agents make mistakes (otherwise we would be in equilibrium). For this very reason economic theorists have had great difficulty in modelling disequilibrium. In order to explain mistakes we need to understand not only the problem faced by agents, but the reasoning of agents, their method of solution. Artificial intelligence provides a practical framework for modelling the reasoning of economic agents in such situations. In strategic environments, agents can actually do better by behaving non-optimally. In such situations, it is thus in agents’ strategic interests to make ‘mistakes’. The actual method of reasoning used to solve the agents’ problems can then be used as a form of precommitment, to influence the eventual outcomes.

This chapter has sought to define the limits and possibilities for artificial intelligence in economic theory, rather than make a positive and substantive contribution and application as found in other papers in Moss and Rae (1992). Whilst I do not see artificial intelligence as a new paradigm that will necessarily replace and

supplant the orthodox economic model of rationality, it clearly has a great potential role, and one that will clearly become very important in future years.

¹ “Compact” is a mathematical term meaning a set is closed and bounded.

² A function $y=f(x)$ maps a value of x onto a single value of y ; a *correspondence* maps a value of x onto one or more values of y .




Figure 7.1 ϵ -optimization.

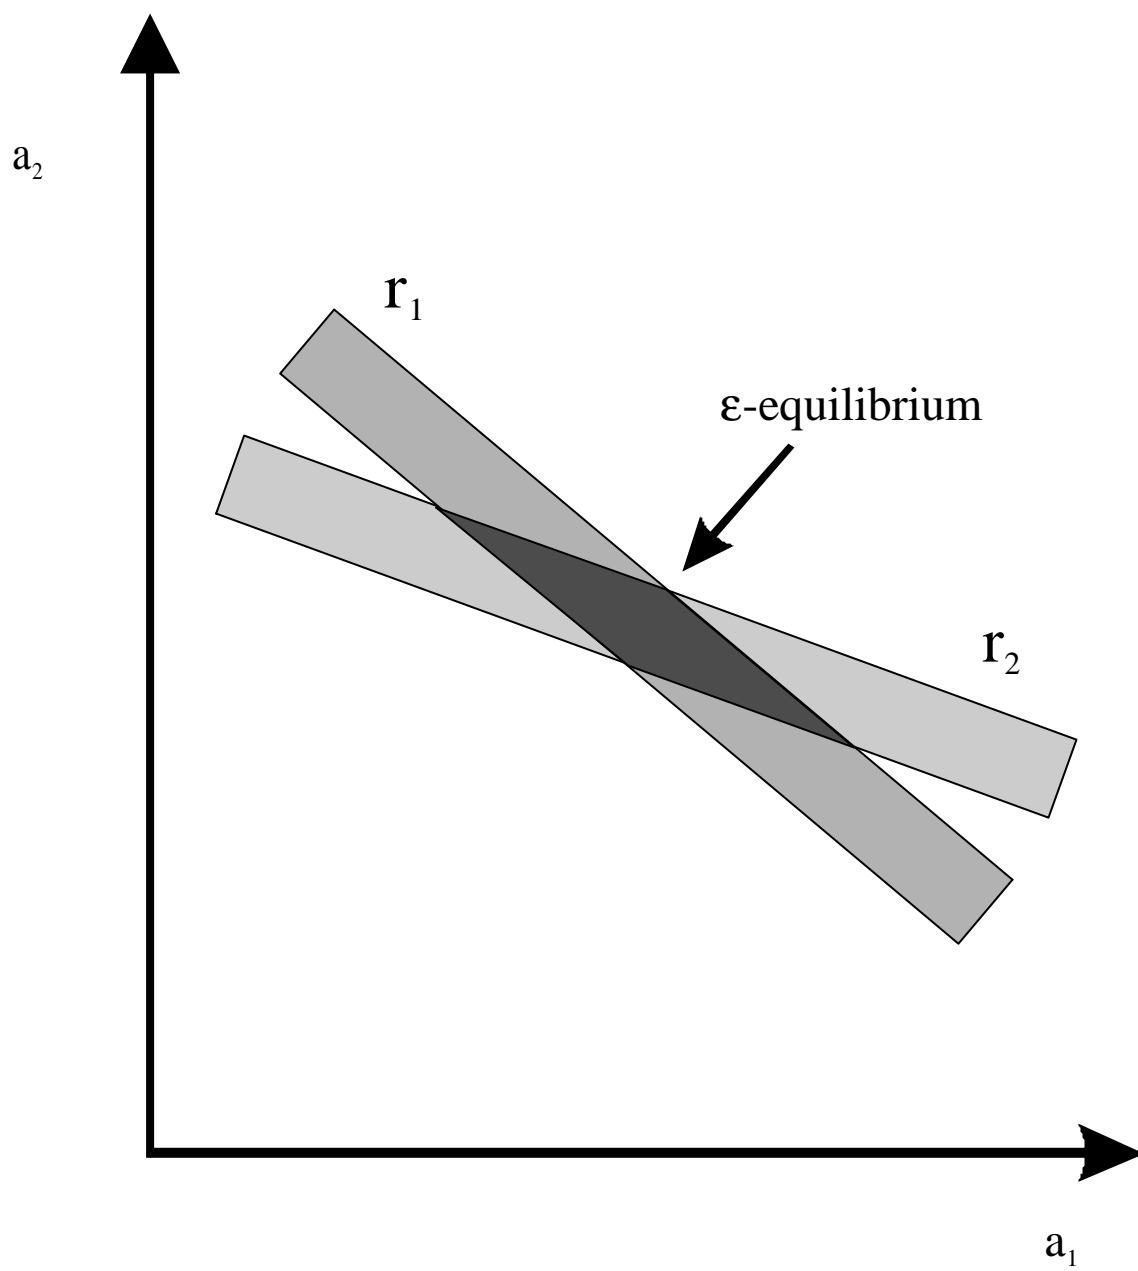
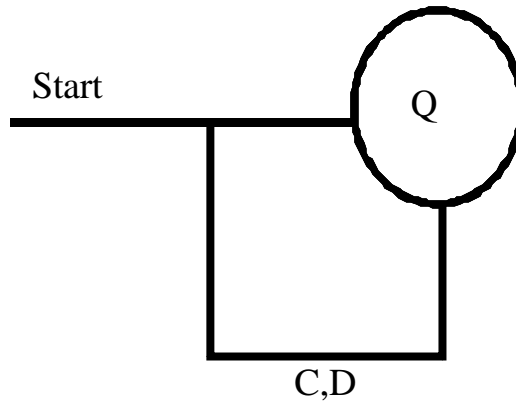



Fig. 7.2 Nash ϵ -equilibria

(a) "Onestate" Machine
 $Q = C \text{ or } D$



(b) "Tit-for-tat" machine



(c) "Grim-Strategy" Machine.




Fig. 7.3 Strategies as finite automata.




Fig. 7.4 Explaining a Mistake




Fig. 7.5 Cournot Duopoly.