30 April, 1998 (saved at 1007)

Modeling R&D Strategy as a Network Search Problem

SCOTT MOSS, BRUCE EDMONDS AND HELEN GAYLARD
{S.MOSS, B.EDMONDS, H.GAYLARD}@MMU.AC.UK
CENTRE FOR POLICY MODELING
MANCHESTER METROPOLITAN UNIVERSITY

1 The setup and its rationale

It is customary in formal models of technological change to treat R&D as a process of directed search. In some analyses, of which the work by Nelson and Winter (1982) is seminal, there is an exogenous development path described by a logistic function reflecting a view that the early stages of a line of technological evolution is difficult but becomes easier as the technology becomes better understood and then more difficult again as the benefits of the core technology are exhausted. Rosenberg (1975, 198x) and Dosi (198x) argue for a connectedness in the development of a technology based in Rosenberg's view on the experience of the innovators and in Dosi's view on innovators' experience and some coherence of the" technological paradigm" on which the stream of innovations are based.

A further set of issues turns on such technological complementarities as in airframe and aeroengine design. Nelson and Winter (1982, pp xxx-xx) have argued that improvements in (say) airframe design to allow airliners to fly higher and faster yield potential benefits that can only be realised once improvements in aeroengine design and materials are available actually to drive the airliner to the higher altitudes and at faster speeds.

One way of thinking about this connectedness within a technological paradigm or core technology is in terms of the diffuseness of the entry points to the technology and the range of developments which follow from each state of technological understanding as well as the complementary technological developments essential to realise a new innovation or its benefits. There is no obvious reason to believe that technological development paths are unique so that any representation of technological possibilities should allow us to express the diversity of development paths leading to a given technology or set of similar technological developments.

Additionally, some developers might give greater emphasis to in-house while others are willing to buy in new developments either by licence or embodied in new plant and equipment. Other characteristics of R&D strategies are important. We know from our work with major UK-based companies that R&D project teams scan a range of possibilities and try to keep more than one alive as they pursue a main line of development in case that line turns out to be a technological *cul de sac*. Moreover, some companies are likely to consider a wider range of possible developments than are others.

Based on these, and other, observations we would expect a model of technological change to have the following characteristics.

1) New technological developments are dependent on the existing technologies in a very specific manner. The knowledge about the general level of technological sophistication is not enough, you also need to know the structure and relatedness of those technologies, i.e. the feasibility of technical innovation is context- and path-dependant.
2) The future uses of technologies to develop other new technologies is largely unpredictable. Firms need to explore the possibilities in some way. For example, they can attempt to develop a large technological base in-house or opportunisticaly look for technologies that can be quickly developed based on existing technologies in the market.
3) There is a difference between technologies that are developed in-house and those that are bought in or licensed, especially when there is a significant gap between the in-house capability and the externally acquired technology.
4) There are at least two type of knowledge that are important to technological innovation: knowledge of the relationship between technologies and the know-how necessary to implement a technology. The point of pilot-studies is primarily to gain knowledge of the former sort.

In order to capture these issues, we will represent the underlying core technology or technological paradigm as a network over which R&D teams search. The network is acyclic on the grounds that experienced R&D teams will not typically discover exactly the same technological advance more than once. Each network node is given a value which is the realization of a random number in the unit interval and represents the productivity increase achieved when an agent acquires that node. Similarly, each link between nodes is given a random value in the [0,10) interval indicating the cost of traversing from a node to the parent node.

We set as parameters the number of nodes in the network, the number of leaf nodes (*i.e.* the number of nodes with links to, but none from, them) and the maximum span of (*i.e.* the maximum number of links from) each node which is not a leaf node. In the models reported here, the number of nodes was 1000 and the maximum span was 5. The number of leaf nodes in each model was different.

The networks are created by the following algorithm:

1) Generate an ordered list of nodes $N = [n_1, n_2, \ldots, n_n]$.

2) Let $N_P = [n_1, n_2,...,n_p]$, where p=n-c and c is the number of leaf nodes
3) Starting with n_1, for each $n_i \in N_P$ in turn
    a) generate a random number r in the [1,5] interval;
    b) select at random r nodes in the list $[n_{i+1}, ni+_2,...,n_n]$ to be the children of n_i;
    c) create a link from n_i to each of its child nodes.

This algorithm gives us a network such as in Figure 1. The leaf nodes of the network, $n_{16},...,n_{20}$ are the possible points of entry into the technology network for R&D teams of the enterprises. All leaf nodes are labelled with 0s to indicate that they are free to acquire but, once acquired, confer no productivity improvements or production cost reductions. they are simply entry points to the network for individual R&D teams. Nodes with children are acquired, and therefore confer the labelled value of productivity improvements, once all of their children have been acquired. So, for example, starting from node n_{16} in Figure 1, node n_6 is then acquired automatically and node n_1 can be *visited* from node n_6. However, to acquire node n_1 in the sense of gaining its implied productivity improvements would require the prior acquisition of node n_3 which could not itself be acquired before the acquisition of nodes n_7 and n_8 (since n_6 must already have been acquired). But, to acquire node n_7, for example, requires the prior acquisition nodes of n_{12} and n_{13} and their children.

In this paper we distinguish between visiting and acquiring a node. When a node is visited, the R&D team discovers the value of the node and the links to and from that node. A node is acquired when it has been visited and all of its children have been acquired. We assume that R&D teams can visit any node which has a link from any node they have already visited. But in any time period these visits can be made only by traversing one link from a previously visited node. In the example given above, there is an easy path from the leaf node n_{16} to node n_1 but to acquire n_1 turns out to be difficult, expensive and time-consuming.

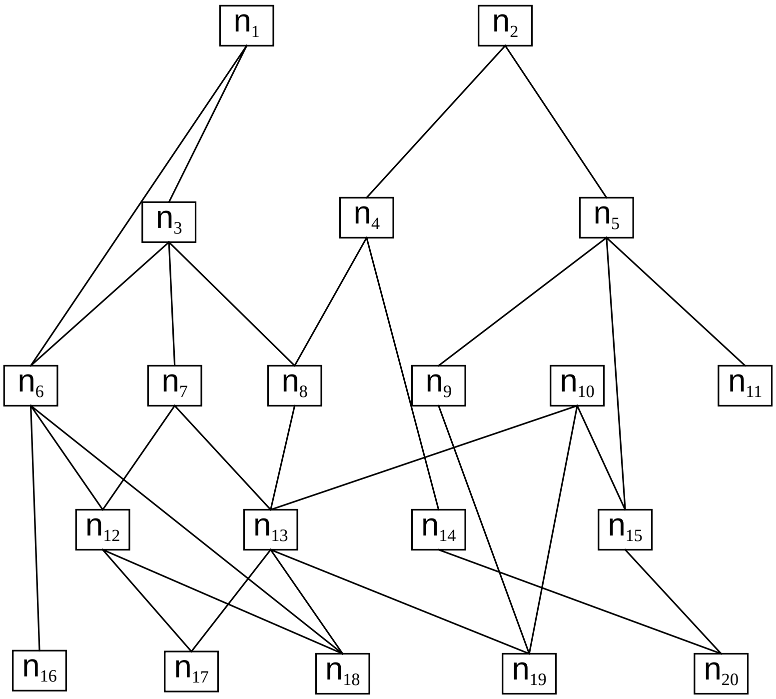

Figure 1: A technology network

2 R&D strategies.

Four R&D search strategies were identified for the models reported here. Agents could engage in a breadth-first search of the technology network or a depth-first search. Additionally and separately, they could restrict their acquisition of nodes to those they had acquired only by search or they could buy nodes that they had visited but not acquired. Each agent was assigned one of the four possible strategy combinations: depth-first and search only, depth-first and a willingness to buy in technology, breadth-first and search only, breadth-first and a willingness to buy in technology. The implementation of these strategies involves the nine rules described below.

1) Node search: breadth-first, search-only.

If all visited nodes have been acquired, then identify the parents of the acquired nodes. Sort those nodes in order of the cost of visiting them from a previously acquired node.

Visit as many of the desired nodes for sale as can be afforded within the technology search budget of the agent.

2) Node search: breadth-first, willing-to-buy, purchase-node.

If all visited nodes have been acquired, then identify which among the parents of the acquired nodes are available for sale by other agents. Sort those which are for sale in order of purchase price. Starting with the cheapest, buy as many of the desired nodes for sale as can be afforded within the technology search budget of the agent.

3) Node search: breadth-first, willing-to-buy, search-for-node.

If all visited nodes have been acquired, then identify which among the parents of the acquired nodes are not available for sale by other agents. Sort those which are not for sale in order of the cost of visiting them from a previously acquired node. Visit as many of the desired nodes for sale as can be afforded within the technology search budget of the agent after deducting the cost of any node purchases.

4) Node search: depth-first, search-only.

If all visited nodes have been acquired, then identify the parents of the acquired nodes and visit the parent node which is the cheapest in the sense that the arc between that node and any acquired node has the label of smallest value among the arcs to any of the acquired nodes.

5) Node search: depth-first, willing-to-buy, purchase-node.

If all visited nodes have been acquired, then identify which among the parents of the acquired nodes are available for sale by other agents. Sort those which are for sale in order of purchase price. Buy the cheapest of the desired nodes for sale that can be afforded within the technology search budget of the agent.

6) Node search: depth-first, willing-to-buy, search-for-node.

If all visited nodes have been acquired and all of the parents of the acquired nodes that are for sale have already been purchased, then sort the parent nodes which are not for sale in order of the cost of visiting them from a previously acquired node. Visit the cheapest of the desired nodes for sale that can be afforded within the technology search budget of the agent after all purchases of nodes.

7) Child node search: search-only.

If any nodes have been visited but not acquired, then sort the children of such nodes in order of the values of the arcs from visited nodes to their children. Starting with the cheapest, visit as many of the desired child nodes as can be afforded within the technology search budget of the agent.

8) Child node search: willing-to-buy, purchase-node.

If any nodes have been visited but not acquired and are for sale by other agents, then sort such nodes in order of the values of the labels on the arcs from visited nodes. Starting with the cheapest, buy as many of the desired nodes for sale as can be afforded within the technology search budget of the agent.

9) Child node search: willing-to-buy, search-for-node.

If the parents of one or more acquired nodes have been visited but not themselves acquired, then sort the unacquired children of such nodes in order of the values of the labels on the arcs from the parents to the children. Starting with the cheapest, visit as many of the desired child nodes as can be afforded within the technology search budget of the agent after all purchases of nodes. (This rule ensures that agents willing to buy in technology will only search one level down from nodes they have identified as being among those they want to move to from their acquired technology. They will not undertake a full development project going back down the technology network. When the node becomes available for purchase, then that option will always be taken.)

3 The simulation experiments

We ran two simulation experiments with this setup. As indicated above, the technology network contained 1000 nodes and the maximum span from each node was 5 in both experiments. In one experiment the number of leaf nodes was 10 and in the other the number of leaf nodes was 100. The idea was to determine whether, in this setup, breadth-first or depth-first search strategies were more efficient and the value of buying in technology when possible instead of relying entirely on in-house development.

In terms of the number of nodes acquired, it is clear from these experiments that breadth-first search was always much the most successful. The results are shown in Figures 2 and 3.

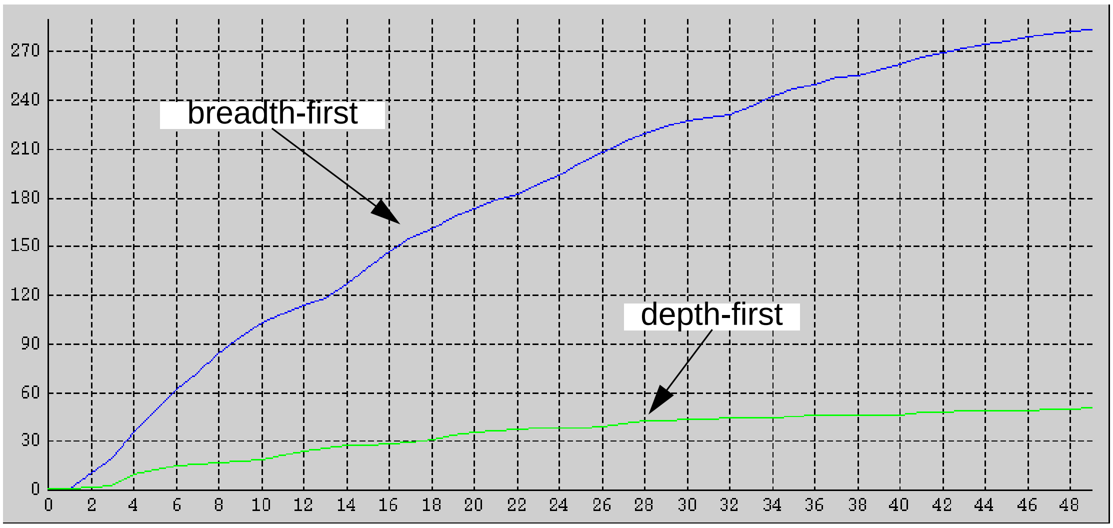

Figure 2: Average number of nodes acquired in 10-leaf-node network

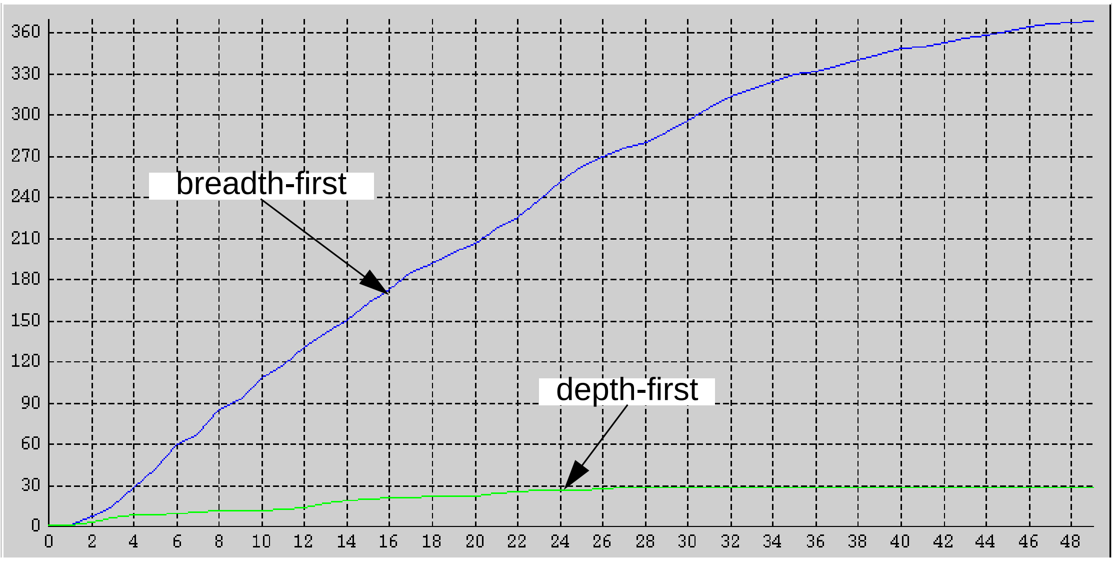

Figure 3: Average number of nodes acquired in 100-leaf-node network

In the representation of the relatively specific technology (10 leaf nodes), the breadth first search was still turning up additional nodes at the end of 50 periods and had reached more than 270 nodes acquired on average among breadth-first searchers. In the same simulation, depth-first searchers had acquired on average about 50 nodes and continued productivity increases were relatively small. In the 100-leaf-node representation, both breadth- and depth-first searches had turned up fewer nodes in the first 50 periods (360+

and 30 nodes, respectively, on average) but the dominance of breadth-first search over depth-first search on the criterion of nodes acquired was much the same.

Certainly one reason for the apparently poor performance of the depth-first search strategy is that it finds root nodes much earlier than does the breadth-first strategy and these offer no scope for further node acquisitions — they are in effect at the limits of the technological regime.

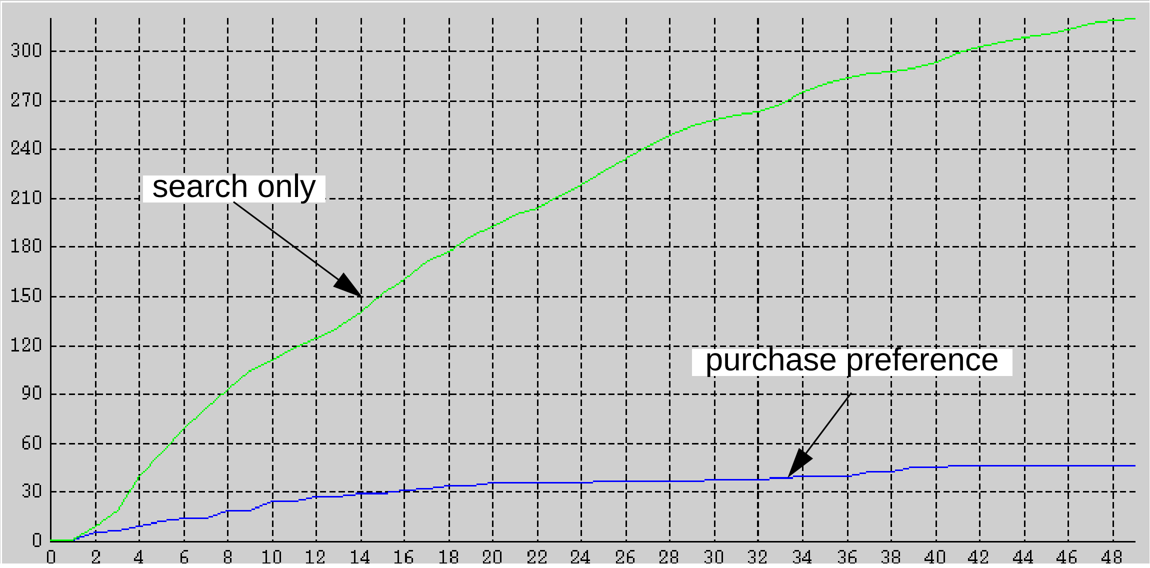

Figure 4: Average number of nodes acquired in 10-leaf-node network

The results reported in Figures 4 and 5 indicate that, in this setup, the largest number of nodes by an order of magnitude or more will be acquired in either technological regime by the agents that search for technology rather than buying it in. This result must stem in some measure from the restriction on search by the agents who will purchase technology when they can — they can only acquire by search the children of nodes for which they have already acquired at least one child.

Of course, measures of success are not limited to number of nodes acquired but also to the financial implications of the strategies. In the setup reported here, the financial position of each agent is determined by net sales revenue and expenditure on technology search (or R&D).

We assumed that each agent sells one unit of output each period. The price of a unit of output is always determined by a 10 per cent mark-up on the average unit cost of production of all agents in the simulation. For each agent, the unit cost of production in any period is $(1+r)^{-1}$ times the unit cost of production in the previous period where r is the sum of the values of the nodes acquired in the previous period. Clearly, agents that acquire fewer nodes will have higher costs of production and therefore less net sales revenue to spend on further node acquisition. In these simulations, we assumed that the technology search budget was always equal to accumulated net revenue less previous expenditures on technology search.

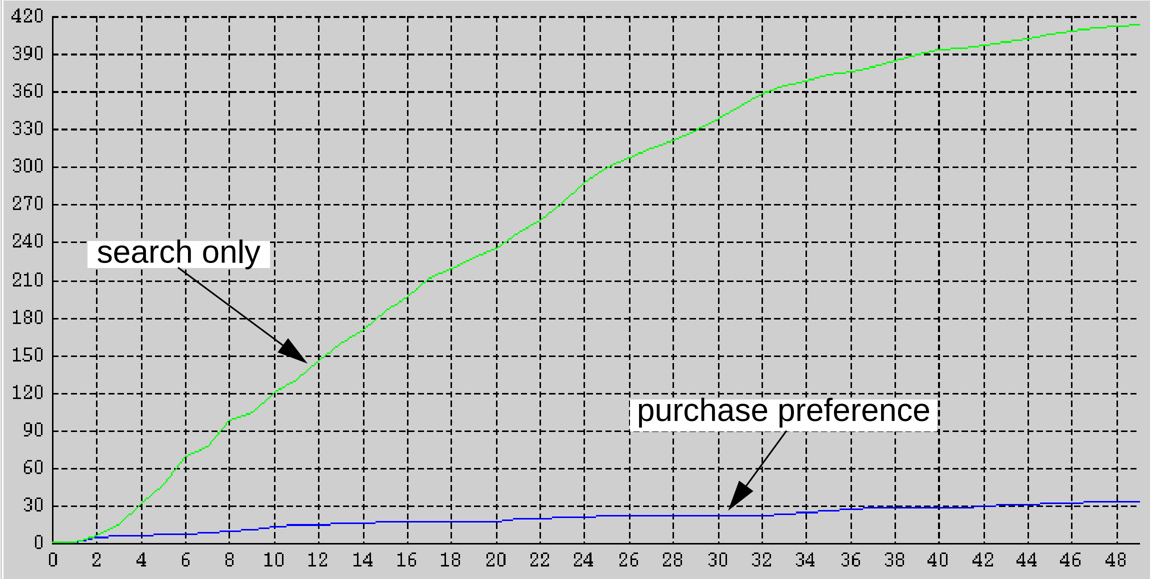

Figure 5: Average number of nodes acquired in 100-leaf-node network

Agents with the highest unit production costs will generate negative net sales revenue. And this is precisely what we observe. On average, agents employing the search-only technology strategy and agents employing the breadth-first search strategy have positive net sales and, therefore, technology search budgets while, after the first few periods, agents buying available technology and engaging in depth-first search have negative technology search budgets so that no search actually takes place. These results are seen in Figures 6 and 7.

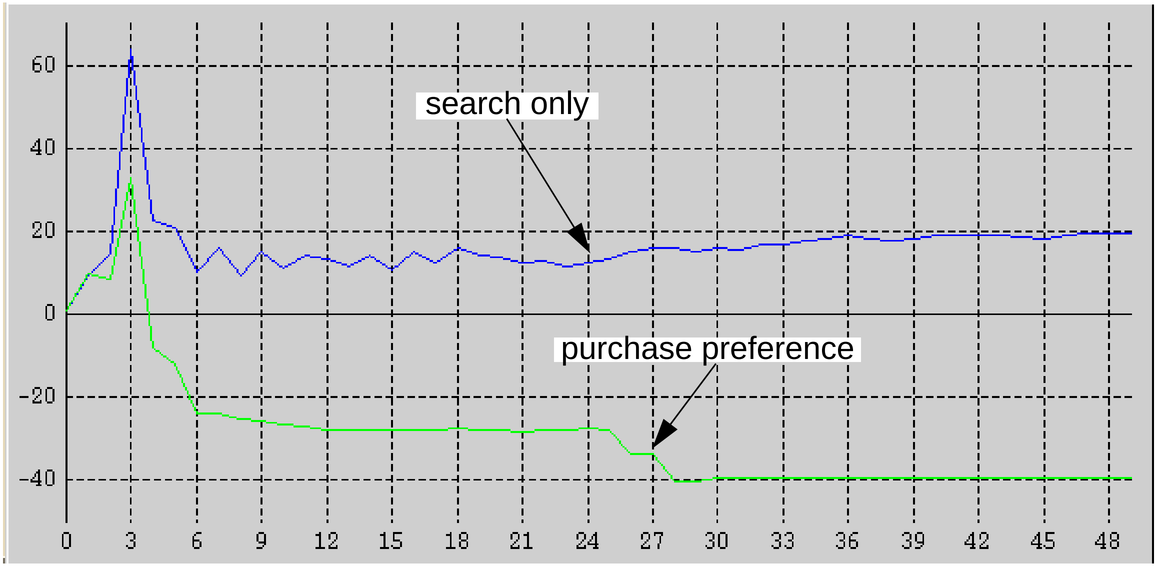

Figure 6: Average cash holdings in 100-leaf-node network

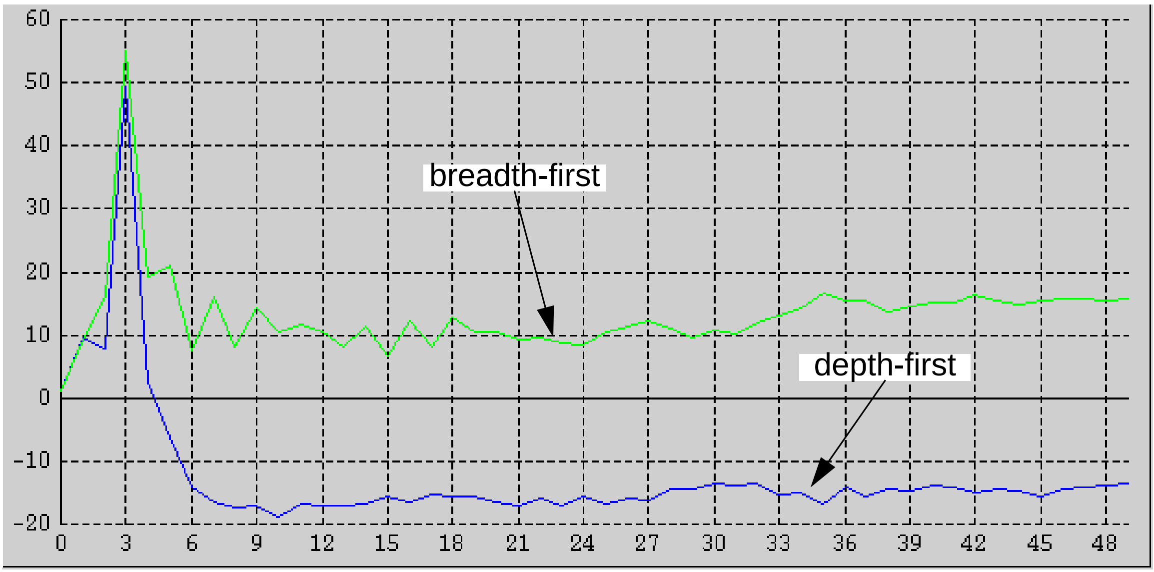

Figure 7: Average cash holdings in 100-leaf-node network

4 Analysis

The results obtained from these simulations are readily explicable as a result of the setup specification.

The principal reason for the relatively poor performance of depth-first search is simply that the algorithm finds root nodes (nodes without parents) and these constitute a *cul de* *sac* from which the agent has no means of escape. A richer model would allow agents to recognize root nodes and give them some strategy for starting their searches from some other node in the network. A larger scale of network might also yield qualitatively different results.

It is generally well recognized by artificial intelligence scientists that a depth-first search is relatively more efficient (and a breadth-first search relatively less efficient) when the search network is deeper and has fewer points of entry. This is what we observe in our simulations since, in the 10-leaf-node network more nodes are acquired by agents adopting depth-first search and fewer nodes are acquired by agents adopting breadth-first search than is found for the 100-leaf-node network.

The greater efficiency of search over technology purchase also follows from the setup specification.

Every node that is acquired reduces the unit production cost of the agent. Agents that only search for technologies must find acquire a larger number of nodes than agents that buy their node-acquisitions. We did not, in these simulations, allow for agents that bought technology to use the resources not allocated to R&D to engage in other activities which might have enhanced their financial positions. We do not, therefore, suggest that our simulation experiments suggest anything very compelling about the virtues of technology licensing *vis a vis* in-house technological development.

5 Developing the models

The purpose of this paper is to demonstrate that network search is a plausible approach to modeling R&D and that different network characteristics can be used to represent the characteristics of individual technological regimes. The networks can themselves be enhanced by, for example, introducing disjuncts so that there are alternative paths to the acquisition of some of the nodes.

The search procedures themselves offer considerable scope for the development of models of R&D teams and how they function within organizations. Our own interests lie in the direction of designing logical formalisms to represent learning and emergent behaviour by R&D teams. Economists are more likely to be interested in the development of path algebras to optimize the search for technological developments. The choice of approach depends on how the analyst conceptualizes R&D. If the scope for technological developments is well understood in advance so that (for example) some kind of stochastic search algorithm is appropriate, then the economist’s approach is appropriate. In cases where the universe of technological possibilities is not already known to the R&D team so that the information which could be found exceeds the information-processing and computational capacities of the agents, then the agents have to make use of heuristics. It is intended that the agents be given a language in which to express, learn and refine such heuristic search strategies. The difference here is the same as that between substantive and procedural rationality.

Our own modeling methodology relies on the declarative modelling techniques. A model is declarative if the current state of the world determines the actions of agents and the ways in which that state will be changed. In other words, declarative models specify reactions to (usually generally specified) states. Economic models, by contrast are typically imperative. They specify a subsequent state on the basis of a given state. In other words, imperative models specify a state to be in rather than an action to take or process to be undergone. Thus, dynamic imperative models yield sequences of states without specifying how they are achieved whereas declarative models specify the actions taken and changes made and it is these which result in subsequent states. In economics, each state is usually an equilibrium whereas in declarative models equilibrium states are the least interesting and are not in any case very likely since they leave no scope for the dynamics inherent in the model specification.

A further advantage of logic-based, declarative modeling techniques is that they enable the modeler to include qualitative considerations without loss of rigour. This will often be important in modeling R&D where the teams will often consider some potential avenues of development to be “promising” or “difficult” without having any numerical measure to represent these views. It is our practice to develop models which rely on qualitative judgements by agents in such a way that simulation outputs include numerical series which can meaningfully be compared with some statistical data series.[1]

6 Conclusion

A model of technological innovation based on network search has been exhibited. This model is credible in that it encodes some aspects of technological context-dependence which is initially unknown to the agents. The model is sufficiently expressive to encapsulate many of the issues of technological innovation and so can be used as a focal point for

---

[1] See Moss and Edmonds (1995) for an example.

a language of discourse on the subject. Further, based on preliminary experiments, the model seems to behave in a way which fits actual observations of such innovation.

7 References

Dosi, G. (1982), “Technological Paradigms and Technological Trajectories”, *Research Policy* **11**:3, pp. 147-162.

Moss, S. and B. Edmonds (1995) “A Formal Preference-State Model with Qualitative Market Judgements”, *Centre for Policy Modeling Report No. 5*, Manchester Metropolitan University (URL: http://www.fmb.mmu.ac.uk/cpm/cpmrep05.html.

Nelson, R. and S. Winter (1982), *An Evolutionary Theory of Economic Change* (Cambridge, MA: Belknap Press).

Rosenberg, N. (1975), *Perspectives on Technology* (Cambridge: Cambridge University Press).

Rosenberg, N. (1982), *Inside the Black Box* (Cambridge: Cambridge University Press).